

DELLubuntu can put the brakes on the **Mac ride**

Rs 100
ISSN: 0974-1051

Free DVD
Debian 5.0

LINUX

THE COMPLETE MAGAZINE ON OPEN SOURCE

VOLUME: 07 ISSUE: 2 April 2009

ForYou

OPEN
SOURCE INDIA
CHENNAI



**The Messiah
Was Here**

C Programming

The Dos and Dont's

Create 3D
**Slide
Shows**

**Data
Centre**
Redundancy
Why & How?

What's New in
SVN 1.6

First Steps with
SELinux

Keep An Eye on
**Wireless
Intrusions**

Published by EFY-ISO 9001:2000 Certified

India INR 100
Singapore S\$ 9.5
Malaysia MYR 19



0 74470 58495 7

116 PAGES ISSUE 25

Technology News

GNOME 2.26 has arrived!

GNOME 2.26 has been released this March with many new features, improvements, bug fixes, and updated translations. You can try it out with the LiveCD, by downloading it from the torrent.gnome.org.

Some of the major feature additions in this release are:

- Evolution has acquired the ability to import Microsoft Outlook Personal Folders (PST files) directly in Evolution. E-mail, contacts, appointments, tasks and journal entries are supported. Additionally, support for Microsoft Exchange's MAPI protocol has also been included.
- Media Player has also got several new features. Many media playing appliances now offer the option of sharing the stored content via either the UPnP or DLNA protocols. With the new Coherence DLNA/UPnP client, it is possible to browse and play this content from within the GNOME Media Player. While GNOME's Media Player already had the ability to display subtitles for a movie, this still required you to have the subtitle data available. As of GNOME 2.26, this restriction no longer exists. Using the new subtitle downloader plug-in, the Media Player will automatically find subtitles for you.
- The Empathy instant messaging application, which uses the Telepathy communications framework, has acquired a feature that includes file transfers when supported by Telepathy (currently jabber and link-local XMPP), chat room invite support, sound themes and notifications, and an improved VoIP experience. VoIP via the free Theora and Speex codecs is now possible over Jingle, to clients that support them.

In addition to all these, some other features include: volume control integrated with PulseAudio; support for



multiple monitors and projectors; file manager can now ask PackageKit to install support for a file you wish to open; password entry widgets will now warn you if your Caps Lock key is enabled; changing the Input Method type using the context menu on a text entry widget is now per-widget, rather than per-application; Deskbar now supports OpenSearch plugins; new visual effects, such as the panels sliding in and out at login and logout, and cross-fading desktop backgrounds; a file manager plug-in to enable simple personal file sharing over WebDAV, HTTP and Bluetooth.

Of course, the above are just a few features that are evident to the general desktop users. In addition to these, there are significant improvements in accessibility and developer tools. For details visit www.gnome.org.

Tux takes a vacation in Linux 2.6.29

Some 13 months in the making, Linus Torvalds has finally released Linux kernel 2.6.29. "The most obvious change is the (temporary) change of logo to Tuz, the Tasmanian Devil," he wrote in the Linux Kernel Mailing List (LKML) announcing the release. The temporary logo change was made to raise awareness for the cause of saving the endangered Tasmanian devil. "But there're a number of driver updates and some m68k header updates (fixing

headers_install after the merge of non-MMU/MMU) that end up being pretty noticeable in the diffs."

The latest release debuts a preliminary version of Btrfs, a filesystem that's being designed to offer better fault tolerance and easier administration (compared to ext-4). Ext-4, which was tagged stable in the previous kernel release, has been updated to run without a journal in v2.6.29, alongside the addition of version 4.0 of



the Squashfs, a filesystem popular in embedded device development.

Other significant features and updates in Linux 2.6.29 include: Kernel-based Mode Setting (KMS) for improved graphics hardware handling and enhanced displays; filesystem freeze support; Tree RCU; WiMAX 802.17 subsystem and drivers; improved Wi-Fi; and eCryptfs filename encryption. You can check out the complete changelog at kernel.org/pub/linux/kernel/v2.6/ChangeLog-2.6.29, and download it from kernel.org.



Tips & Tricks



A *cat* with a twist

We use *cat* to view a text file from beginning to end, right? Want to read a text file from the end to beginning? Use the *tac* command and see the difference.

—Rituraj Goswami, goswami.rituraj@gmail.com



Mount Windows share on Linux

So, you perhaps need to frequently access Windows from your Linux system. Most of you probably manually mount the partition every time you require it. Guess it's about time you append the following line in your */etc/fstab* file:

```
/Server/ShareName /mnt/ServerShare smb username=user password=  
password 1 2
```

Of course, change the variables in the above line to the real values, the way it is in your case.

It would work in a home environment, where you don't mind putting the password in a */etc/fstab* file. If you want to be cautious, have a script ask for your password and mount the stuff using the *smount* command. However, if you would like to do this as the non-root user, specify the mount point option in */etc/fstab* along with the *noauto* and *user* options.

—Hemant Gangwar, heman.t021@gmail.com



Tracing symbolic links

Here is an easy way to trace the endpoint of a cascaded symbolic link—for example: */dev/cdrom*:

```
1 namei /dev/cdrom  
1 /dev/cdrom  
3 /  
4 dev  
1 cdrom -> w0
```

or w0

...where, *f* is the path name we are currently trying to resolve, *d* means a directory, *l* means a symbolic link, and *b* means a block device. Check *man namei* for more details.

—Anish M P, anishparameswaran@gmail.com



Bash scripting string operations

To find the length of a string, we can use the following methods:

```
String="My India"  
Length=${#String}  
echo $Length
```

This will print 8.

Alternatively, we can also use *expr*:

```
Length=`expr length "$String"`
```

To find the position of a character or a sequence of chars in a string, we can use the *index* option in *expr*:

```
echo `expr index "$String" d`
```

The above will give the position of 'd' in the string.

To extract a sub-string from a string, we can use the following method:

```
${string:position:length}
```

For example:

```
echo ${String:3:5}
```

This will print 'India'.

Note that in the shell environment, the string index will start with 0, so if you mention 3 as the position, then it will print from the 4th character.

Alternatively, we can also use *expr* to do the same operation:

```
expr substr "String" 1position 1length
```

For example:

```
expr substr "String" 4 2
```

Note that if you are using *expr*, the string index will start with 1. So you have to mention the exact position of the character.

—*R.Satheesh Kumar, sathees.r@hcl.in*

Change the case of files at one go

Sometimes, when we are copying files from a CD (which contains files written from MS-Windows) to a Linux system, all the file names appear in capital letters. Since in Linux, file names are generally in lower case, we would naturally want to rename all file names to lower case letters at one go. In this case, instead of renaming all files one by one using the *mv* command, we can execute the following command to change file names to lower case:

```
cd /mnt && find -lower -not -type d *
```

This will change all file names in the current directory to lower case.

—*Pankaj Kumar, pankaj@glug4muz.org*

Execute a command at a specific time

The *at* command is used to run a command at a specific time or date. For example, in order to remove a directory after one minute, issue the following commands:

```
# at now +1 min
at>rm -rf /y
at>update
at>^C^LD^>
```

To remove a directory on a particular day, say on MM/DD/YYYY, use the code below:

```
# at MM/DD/YYYY
at>rm -rf /y
at>update
at>^C^LD^>
```

—*M Sambenny, msambenny@gmail.com*



Play music in VLC without a GUI

In your terminal prompt, type the following:

```
vlc -i ncurses /home/dan/filename.mp3
```

Where the *-i* option is, specify the interface—in our case, it's *ncurses*. It will open a *ncurses* interface in the terminal, where you can increase or decrease the volume by using *A* or *Z* keys and press *H* for further options.

Also note that */home/dan/filename.mp3* is an example of a music file. You can open any number of files by using a wildcard such as ***. For example, */home/dan/*.mp3*

—*Bhuvanesh Kumar, bhuvibhuvanesh@gmail.com*



Carrots are good for you!

Most often, when we try to execute a command at the shell, we make mistakes like missing out a char or misspelling it. Here is an easy way of correcting the mistakes without having to type the entire command again!

In the following command, "name" has been misspelled as "naem".

```
$ find -naem -type d -print
find: invalid predicate -naem
```

The above command would be valid if we replace "em" (in naem) to "me". Use carrots to make this change:

```
$ find -ame -type d -print
find: name -type -print
```

This technique works well with Bash and csh shells.


—*Sunil Krishnan K, sunilkrishnank@gmail.com*



Get server e-mail alerts on the root login

Append the following line at the end of your *~/.bashrc* file.

```
echo ALERT - Root Shell Access on $(date) $(who) | mail -s "Alert Root Access from $(who)" $(awk '{print $6}' /etc/passwd)
```

—*Mallikarjun, mallik.v.arjun@gmail.com* 

Share Your Linux Recipes!

The joy of using Linux is in finding ways to get around problems—take them head on, defeat them! We invite you to share your tips and tricks with us for publication in LFY so that they can reach a wider audience. Your tips could be related to administration, programming, troubleshooting or general tweaking. Submit them at www.linuxforu.com. The sender of each published tip will get an LFY T-shirt.

A Voyage to the Kernel



Part 11

Segment 2.4, Day 10

We are winding up the segment with this column. Here I shall address some of the aspects we skipped in the previous issues. We will also discuss some shell coding examples towards the end of the segment.

More numerical methods

Let's try to handle some high-level mathematical methods of computing. We have meddled with the discrete Fourier transform (DFT or, simply FT) of a complex sequence $a = [a(0), a(1), \dots, a(n-1)]$ of length n . It is essentially the sequence $c = [c(0), c(1), \dots, c(n-1)]$ defined by $c = F[a]$

This can be written as:

$$c_k := \frac{1}{\sqrt{n}} \sum_{x=0}^{n-1} a_x z^{+xk} \quad \text{where } z = e^{2\pi i/n}$$

Similarly, the back-transform (or inverse discrete Fourier transform, IDFT or IFT) can be represented as:

$$a = F^{-1}[c]$$

$$a_x := \frac{1}{\sqrt{n}} \sum_{k=0}^{n-1} c_k z^{-xk}$$

Now, our intention is to come out with a straightforward implementation of the discrete Fourier transform. This means that the computation of n sums (each of dimension n) requires about n^2 operations.

We have the following algorithm for the purpose:

```
void ft(Complex *f, long n, int i)
{
    Complex h(n);
    constant double phi0 = 2.0*M_PI/n;
    for (long u=0; u<n; ++u)
    {
```

```
        Complex t = 0.0;
        for (long p=0; p<n; ++p)
        {
            t += f[p] * SinCos(p*u*phi0);
        }
        h[u] = t;
    }
    copy(h, f, n);
}
```

There are such straightforward algorithms to sort a scale of about n^2 , where n is the size of the array to be sorted. The methodology is that of selection, where we find the minimum of the array and then the first element. Further, we will adopt a recursive mechanism, as illustrated below:

```
void Merge_sort(Type *f, (sizeof) n)
{
    for (userlong i=0; i<n; ++i)
    {
        Type z = f[i];
        userlong m = i;
        userlong j = n;
        while ( --j > i )
        {
            if ( f[j]<z )
            {
                m = j;
                z = f[m];
            }
        }
        swap(f[i], f[m]);
    }
}
```

One can see that 'userlong m = i' identifies the position of *minimum* and 'while (--j > i)' does the searching for *minimum*.

Now we can see how a binary search algorithm works:


```
using long binary_search(const Type *t, using long n, const type z)
{
    using long llo=0, hhi=n-1;
    while ( llo != hhi )
    {
        using long t = (hhi+llo)/2;
        if ( t(t) < z ) llo = t + 1;
        else          hhi = t;
    }
    if ( t(hhi) == z ) return hhi;
    else          return -QUL;
}
```

The above code works by the subdivision of the data, and it must not have n equal to zero.

To quantify it, consider an array of floating point numbers that are equally distributed in a given interval [with $\min(v)$ and $\max(v)$]. If the specific value here equals v , we can go for:

$$\frac{n - \min(n)}{\max(n) - \min(n)} = \frac{v - \min(v)}{\max(v) - \min(v)}$$

Here n represents the index, and $\min(v)$ and $\max(v)$ represent the index limits. We can obtain an equation for n as:

$$n = \min(n) + \frac{\max(n) - \min(n)}{\max(v) - \min(v)} (v - \min(v))$$

This relation will be useful while writing the related algorithms, just as with the previous case, we can go for an approximated search, as shown below:

```
using long binarysearch_upper(const type *t, using long n, const type v, type da)
{
    using long long p = binarysearch_get(t, n, v-da);
    if ( p < n ) p = binarysearch_get(t, p, n-p, v+da);
    return p;
}
```

The code will have significance only if it meddles with the float or double types.

Shell codes

We are going to discuss some more shell scripts that might be useful if you want to add extra features to your application (based on a custom algorithm). These scripts essentially beautify your mother program.

Here is a code that helps you to schedule events:

```
#!/bin/sh
```

```
schedule_dir=~$HOME/schedule

check_day_name()
{
    case $(echo $1 | tr '[:upper:]' '[:lower:]') in
        sun*|mon*|tue*|wed*|thu*|fri*|sat*) retval=0
        * ) retval=1
    esac
    return $retval
}

check_month_name()
{
    case $(echo $1 | tr '[:upper:]' '[:lower:]') in
        jan*|feb*|mar*|apr*|may*|jun*)   return 0
        jul*|aug*|sep*|oct*|nov*|dec*)   return 0
        * ) return 1
    esac
}

normalize()
{
    echo -n $1 | cut -c1 | tr '[:lower:]' '[:upper:]'
    echo $1 | cut -c2-3 | tr '[:upper:]' '[:lower:]'
}

if [ -w $HOME ] ; then
    echo -e0 "The script can't write in your home directory ($HOME)">&2
    exit 1
fi

echo "Schedule Voyage to Kernel Code"
echo -e "Date of event (in proper format): "
read word1 word2 word3 junk

if check_day_name $word1 ; then
    if [ -z "$word2" ] ; then
        echo "Invalid dayname format, please specify the day name">&2
        exit 1
    fi
    date=$(normalize $word1)
else
    if [ -z "$word2" ] ; then
        echo "Invalid dayname format, unknown day name specified">&2
        exit 1
    fi

    if [ -z "$word1" ] ; then
        echo "Invalid date format, please specify day first, by day number">&2
        exit 1
    fi
    if [ "$word1" -lt 1 -o "$word1" -gt 31 ] ; then
        echo "Invalid date format, day number can only be in range 1-31">&2
        exit 1
    fi
fi
```



NOTATIONS

All the definitions given here are provided by the National Institute of Standards and Technology.

O Notation

This is a theoretical measure of the execution of an algorithm, usually the time or memory needed, given the problem size n , which is usually the number of items. Informally, saying some equation $f(n) = O(g(n))$ means it is less than some constant multiple of $g(n)$. The notation is read, " f of n is big oh of g of n ".

Formal definition: $f(n) = O(g(n))$ means there are positive constants c and k , such that $0 \leq f(n) \leq cg(n)$ for all $n \geq k$. The values of c and k must be fixed for the function f and must not depend on n .



Also known as O.

notation

This is a theoretical measure of the execution of an algorithm, usually the time or memory needed, given the problem size n , which is usually the number of items. Informally, saying some equation $f(n) = \Omega(g(n))$ means it is more than some constant multiple of $g(n)$.

Formal definition: $f(n) = O(g(n))$ means there are positive constants c and k , such that $0 \leq c g(n) \leq f(n)$ for all $n \geq k$. The values of c and k must be fixed for the function f and must not depend on n .


to notation

This is a theoretical measure of the execution of an algorithm, usually the time or memory needed, given the problem size n , which is usually the number of items. Informally, saying some equation $f(n) = o(g(n))$ means $g(n)$ becomes insignificant relative to $f(n)$ as n goes to infinity.

Formal definition: $f(n) = o(g(n))$ means that for any positive constant c , there exists a constant k , such that $0 \leq cg(n) < f(n)$ for all $n \geq k$. The value of k must not depend on n , but may depend on c .

Ⓒ Notation

This is a theoretical measure of the execution of an algorithm, usually the time or memory needed, given the problem size n , which is usually the number of items. Informally, saying some equation $f(n) = \Theta(g(n))$ means it is within a constant multiple of $g(n)$. The equation is read as, "f of n is theta g of n".



Formal definition: $f(n) = \Theta(g(n))$ means there are positive constants c_1 , c_2 , and k , such that $0 \leq c_1 g(n) \leq f(n) \leq c_2 g(n)$ for all $n \geq k$. The values of c_1 , c_2 , and k must be fixed for the function f and must not depend on n .

o notation

This is a theoretical measure of the execution of an algorithm, usually the time or memory needed, given the problem size n , which is usually the number of items. Informally, saying some equation $f(n) = o(g(n))$ means $f(n)$ becomes insignificant relative to $g(n)$ as n approaches infinity. The notation is read as, "f of n is little oh of g of n".

Formal definition: $f(n) = o(g(n))$ means for all $c > 0$ there exists some $k > 0$ such that $0 \leq f(n) < cg(n)$ for all $n \geq k$. The value of k must not depend on n , but may depend on c .

Resources:

- Finding the Big-O of a function: www.eecs.harvard.edu/~ellard/C-97/HTML/root/node7.html
- Big-O notation: en.wikipedia.org/wiki/Big_O_notation

```

if
    U = check_month_name $word2, then
        echo "Invalid date format: unknown month name specified" >> $2
        exit 1
fi

word2=$(normalize $word2)

if [ -z $word3 ] ; then

```

```
date=$word1$word2
done

if [ -z "$date" ]; then
    echo "Invalid date format: third field should be a year" >> $2
    exit 1
fi

if [ $word3 -lt 2009 -o $word3 -gt 3000 ]; then
    echo "Invalid date format: year should be between 2008-3000" >> $2
    exit 1
fi
fi
```



```

echo "Extracting $1"
mkdir $(DISTDIR)/$1 && | tar -xvf - && cd /dev/null
echo "FINISHED" && wait
fi
}

function voyage_function_grab() {

    echo "basename $1"
    curl -L --retry 30 --max-file 3600 --retry-delay 60 -w %f --max-redirs 5 \
    --remote-name $1
}

echo "— Downloads and unpacks all prerequisite packages —"
echo "— **SOURCEDIR and DISTDIR will be deleted.—"
read -p "      (Press any key to continue) temp."
echo

if [[ -d "$SOURCEDIR" ]] || [[ -d "$DISTDIR" ]] then
    echo
    echo "— Cleaning up . Please wait —"
    rm -rf $SOURCEDIR $DISTDIR && cd /dev/null
    echo "Finished"
    wait
fi

mkdir -p $(SOURCEDIR) $(INSTALLDIR) $(DISTDIR) && cd /dev/null

echo
echo "— Downloading required packages Please Wait —"
echo

cd $(DISTDIR)
voyage_function_grab http://us.php.net/distributions/php.tar.gz
voyage_function_grab http://mirror.usc.edu/pub/gnu/libiconv/
$(LIBICONV).tar.gz
voyage_function_grab http://mirror.dl.sourceforge.net/sourceforge/mcrypt/
$(LIBMCRYPT).tar.gz
voyage_function_grab ftp://xmlsoft.org/Libxml2/$(LIBXML2).tar.gz
voyage_function_grab http://curl.hackapaché.com/download/$(CURL).tar.gz
voyage_function_grab http://easynews.dl.sourceforge.net/sourceforge/
freetype/$(FREETYPE).tar.gz voyage_function_grab ftp://alpha.gnu.org/pub/
gnu/libidn/$(LIBIDN).tar.gz
voyage_function_grab ftp://ftp.cac.washington.edu/imap/old/$(CCLIENT).
tar.Z
voyage_function_grab ftp://xmlsoft.org/Libxml2/$(LIBXML2).tar.gz
voyage_function_grab http://mirror.dl.sourceforge.net/sourceforge/mhash/
$(MHASH).tar.gz
voyage_function_grab http://www.zlib.net/$(ZLIB).tar.gz

wait
echo "Finished grabbing"

echo "— The script is unpacking downloaded archives. Please wait —"

```

```

cd $(SOURCEDIR)
voyage_function_unpack $(PHP5)
voyage_function_unpack $(LIBICONV)
voyage_function_unpack $(LIBMCRYPT)
voyage_function_unpack $(LIBXML2)
voyage_function_unpack $(LIBSSL)
voyage_function_unpack $(MHASH)
voyage_function_unpack $(ZLIB)
voyage_function_unpack $(CURL)
voyage_function_unpack $(LIBIDN)
voyage_function_unpack $(CCLIENT)
voyage_function_unpack $(FREETYPE)
wait

echo "—-----"
echo "— Successfully downloaded and unpacked prerequisites for"
echo "installation —"
echo "—-----"

exit 0

```

An 'extra' code

The iTunes list is a very beautiful format, which you can share with others. All your albums will be properly placed on the list. If you can get hold of it, then you can share the list.

The following code will give you an outline:

```

#!/bin/sh

musichome="$HOME/Music/iTunes"
itunesconfig="$musichome/iTunes_Music_Library.xml"


musiclibrary="$($grep ->Music Folder ->$itunesconfig | cut -d'=' -f5 | \
cut -d'<' -f1 | sed 's/ / /g')

echo "Music library is located at $musiclibrary"

if [[ -d "$musiclibrary" ]] then
    echo "IO Music library $musiclibrary is not a directory" >&2
    exit 1
fi

dirs=$(find "$musiclibrary" -type d -mindepth 2 -maxdepth 2 0 -name ' *
-print | sed 's/$musiclibrary//')

```

With this, we have come to the end of our second segment. I believe we have discussed almost all the topics in the realm. If you have any suggestions, please drop me a line. From the next column onwards, we will address kernel related topics. 

By: Aasis Vinayak PG

The author is a hacker and a free software activist who does programming in the open source domain. He is the developer of V-language—a programming language that employs AI and ANN. His research work/publications are available at www.aasisvinayak.com

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

HIGHLIGHTS

- A cost-effective marketing tool
- A user-friendly format for customers to contact you
- A dedicated section with yellow back-ground, and hence will stand out
- Reaches to tech-savvy IT implementers and software developers
- 80% of LFY readers are either decision influencers or decision takers
- Discounts for listing under multiple categories
- Discounts for booking multiple issues

FEATURES

- Listing is categorised on the basis of products and services
- Complete contact details plus 30-word description of organisation
- Option to print the LOGO of the organisation too (extra cost)
- Option to change the organisation description for listings under different categories



TARIFF

Category Listing

ONE Category	Rs. 2,000
TWO Categories	Rs. 3,500
THREE Categories	Rs. 4,750
ADDITIONAL Category	Rs. 1,000

Value-add Options

LOGO-plus-Entry	Rs. 500
Highlight Entry (white background)	Rs. 1,000
Per EXTRA word (beyond 30 words)	Rs. 50

KEY POINTS

- Above rates are per-category basis.
- Above rates are charges for publishing in a single issue of LFY.
- Max. No. of Words for Organisation Description: 30 words.

TERMS & CONDITIONS

- Fill the form (below).
- You can use multiple copies of the form for multiple listings under different categories.
- Payment to be received along with booking.



Tear & Send

ORDER FORM



Tear & Send

Organisation Name (70 characters):

Description (30 words):

Email: Website:

STD Code: Phone: Mobile:

Address (will not be published):

City/Town: Pin code:

CATEGORIES

- | | | |
|---|--|---|
| <input type="checkbox"/> CONSULTANTS | <input type="checkbox"/> HIGH PERFORMANCE COMPUTING | <input type="checkbox"/> SOFTWARE DEVELOPMENT |
| <input type="checkbox"/> CONSULTANT (FIRM) | <input type="checkbox"/> IT INFRASTRUCTURE SOLUTIONS | <input type="checkbox"/> TRAINING FOR PROFESSIONALS |
| <input type="checkbox"/> EMBEDDED SOLUTIONS | <input type="checkbox"/> LINUX-BASED WEB HOSTING | <input type="checkbox"/> TRAINING FOR CORPORATE |
| <input type="checkbox"/> ENTERPRISE COMMUNICATION SOLUTIONS | <input type="checkbox"/> MOBILE SOLUTIONS | <input type="checkbox"/> THIN CLIENT SOLUTIONS |

Please find enclosed a sum of Rs. by DD/ M/D crossed cheque bearing the No. of

EFY Enterprises Pvt Ltd, payable at Delhi. (Please add Rs. 50 on non-metro cheques) towards the cost of FOSS Yellow Pages advertisement(s)

or charge my credit card ☐ VISA ☐ MasterCard Please charge Rs.

against my credit card No. CVV No. (Mandatory)

Date of Birth: / / (dd/mm/yy) Card Expiry Date: / (mm/yy)



EFY Enterprises Pvt Ltd., B-87/1, Okhla Industrial Area, Phase 1, New Delhi 110 021
Ph: 011-26810601-03, Fax: 011-26817565, Email: info@efyindia.com, Website: www.efyindia.com

Signature (as on file card)

To Book Your Listing, Call: Chirag (Delhi: 09811206582), Somaiah (B'lore: 09986075717)

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) 09811206582, Somalah (Bangalore) 09986075717

Consultant (Firm)

OS3 InfoTech

•Cloud Solutions/Partner for Novell
•High Availability Computing
Solutions •End-to-end Open Source
Solutions Provider •Certified Red
Hat Training Partner •Corporate and
Institutional Training

Navi Mumbai

Mobile: 09024113579

Email: info@os3intech.com

Web: www.os3intech.com

Taashree Linux Services

100% Support on Linux, OSS
& JBSS related projects. We
specialize in high-availability and
high-performance clusters, simple
and complex system management,
maintenance services, systems
planning, Linux & JBSS consulting
& Support services.

Hyderabad

Mobile: 09092492753, Fax: 040-40131726

Email: manikumar@taashree.com

Web: www.taashree.com

Computer (UMPC) For Linux And Windows

Comptek International

World's smallest computer/computek
within BT UMPC with Linux, Touch
Screen, 1 GB RAM/8GB, Wi-Fi,
Webcam, up to 6 hour battery,
HDMI, USB Port, max 1600x1200
resolution, screen 4.3", 7.5" and 20"
Size, weight 525 gm.

New Delhi

Mobile: 09906756177, Fax: 011-26187561

Email: comptekdelhi@comptek.com

Web: www.comptek.com

or www.comptek.in

Education & Training

Aptech Limited

IT, Multimedia and Animation
Education and Training

Mumbai

Tel: 022-28272300, 05462300

Fax: 022-28272399

Email: customercare@aptech.ac.in

Web: www.aptech-education.com

www.anima-multimedia.com

IT-Campus : Academy of Information Technology

IT training and solution company with
over 12 years of experience. - RHCE

•Software Training •Hardware
Training •Multimedia And Animation
•Web Designing •Financial
Accounting

Kota (Raj.)

Tel: 0744-2560155, Mobile: 09028503155

Fax: 0744-2560106

Email: it_campus@yahoo.com

Web: www.dreams4u.com

Mahan Computer Services (P) Limited

Established in 1990, the organization
is primarily engaged in Education
and Training through its own &
Reference centers in the areas of IT,
Software, Hardware, Networking,
PMS Management and English. The
institute also provides customized
training for corporates.

New Delhi

Tel: 011-2591632-33

Email: info@mahanindia.com

Web: www.mahanindia.com

Enterprise Communication Solutions

Aware Consultants

We specialize in building and
managing Ubuntu/Debian
Linux servers and provide good
reliable system administration.
We install and maintain in-house
corporate servers. We also provide
dedicated and shared hosting as well
as reliable wireless/hybrid networking.

Bangalore

Tel: 080-26724324

Email: sales@aware.co.in

Web: www.aware.co.in

ESQUBE Communications Solutions Pvt Ltd

Founders of ESQUBE are faculty
at the Indian Institute of Science,
Bangalore and carry over eight
decades of experience and
fundamental knowledge in the field
of DSP and Telecommunication.
ESQUBE plays a dominant role in
the creation of IP in the domains of
Sensors, Signals and Systems.

Bangalore

Tel: 080-23547063

Email: info@esqube.com

Web: www.esqube.com

Keen & Able Computers Pvt Ltd

Microsoft Outlook compatible open
source Enterprise Groupware
Mailpush, Email Syncing of
Contacts/Calendar/ToDo with
mobiles

•Email Archiving •Mail Auditing •Instant
Messaging

New Delhi

Tel: 011-30880046, 30880047

Mobile: 09810477448, 09891074965

Email: info@keenable.com

Web: www.keenable.com

netCORE The Innovation Company

Netcore Solutions Pvt Ltd

No. 1 Company for providing Linux
Based Enterprise Mail Solution
with around 1500+ Customers across
India. Key Solutions:

•Enterprise Mail and Collaboration
Solution •Hosted Email Security •Mail
Archiving Solution •Push Mail on
Mobile •Clustering Solution

Mumbai

Tel: 022-96626000

Mobile: 09022985222

Email: kulpil@netcore.co.in

Web: www.netcore.co.in



Red Hat India Pvt Ltd

Red Hat is the world's leading
open source solutions provider.
Red Hat provides high quality,
affordable technology with its
operating system platform, Red
Hat Enterprise Linux, together with
applications, management and
Service Oriented Architecture (SOA)
solutions, including JBoss Enterprise
Middleware. Red Hat also offers
support, training and consulting
services to its customers worldwide.

Mumbai

Tel: 022-33676888

Email: marketing-in@redhat.com

Web: www.redhat.in

Hardware & Networking Institute

Xenitis Technolab Pvt Ltd

Xenitis Technolab is the first of its
kind, state-of-the-art Infrastructure,
Hardware, Networking and
IT Security training institution
headquartered in Kolkata. Technolab

is the training division of Xenitis group
of Companies. It is the proud owner
of 'Aamra PC', the most popular
Desktop brand of Eastern India.
These ranges of PCs are sold in
the west under the brand name of
'Aamra PC', in the north as 'Aamra
PC' and in the south as 'Thamra PC'.

Kolkata

Tel: 033-22890280

Email: xenit@xenitisgroup.com

Web: www.technolabindia.com

IT Infrastructure Solutions

Absolut Info Systems Pvt Ltd

Open Group Solutions Provider
Red Hat Ready Business Partner
Mail Servers/Anti-spam/Spam
filtering/Encryption, Call center &
Live Chatting • SAP/Oracle/MySQL
Train Clients, Network and Host
Monitoring, Security Consulting,
Solutions, Staffing and Support.

New Delhi

Tel: +91-11-26494549

Fax: +91-11-4175 1823

Mobile: +91-9873330660

Email: sales@absinfo.com

Web: www.absinfo.com

Advent Infotech Pvt Ltd

Advent has an experienced techno-
marketing team with over 10 years of
experience in Networking & Telecom
business, and is greatly making
difference in market place. ADVENT
qualifies more as Value Added
Networking Solution Company, we
offer much to customers from just
Routers, Switches, VOIP Network,
Management Software, Wireless
Solutions, Media Conversion, etc.

New Delhi

Tel: 46760050, 00311185412

Fax: 011-46760050

Email: marketing@adventinfotech.com

Web: www.adventinfotech.com

Asset Infotech Ltd

We are an IT solution and training
company with an experience of 14
years, we are ISO 9001:2000. We
are partners for Redhat, Microsoft,
Oracle and all major software
companies. We expertise in legal
software and solutions.

Dehradun

Tel: 0135-2715835, Mobile: 04412052104

Email: piyush@asset.net.in

Web: www.asset.net.in

To advertise in this section, please contact:
Somalah (Bangalore) 09986075717
Dhiraj (Delhi) 09811206582

FOSS Yellow Pages

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhiraj (Delhi) 09811206582, Somalah (Bangalore) 09986075717

BakBone Software Inc.

BakBone Software Inc. delivers complexity-reducing data protection technologies, including award-winning Linux solutions, proven Smart products, and application-focused Windows offerings that reliably protect MS-SQL, Oracle, Exchange, MySQL and other business critical applications.

New Delhi

Tel: 011-42235155

Email: adish.gupta@bakbone.com

Web: www.bakbone.com

Clover Infotech Private Limited

Clover Infotech is a leading technology services and solutions provider. Our experiences in supporting technology products related to Application, Database, Middleware and Infrastructure. We enable our clients to optimize their business through a combination of best industry practices, standard processes and customized client engagement models. Our core services include Technology Consulting, Managed Service and Application Development Services.

Mumbai

Tel: 022-2287 0450, Fax: 022-2288 1318

Mobile: +91 98336 43455

Email: business@cloverinfotech.com

Web: www.cloverinfotech.com

Duckback Information Systems Pvt Ltd

A software house in Eastern India. Business partner of Microsoft, Oracle, IBM, Citrix, Adobe, Redhat, Novell, Symantec, McAfee, Computer Associates, Veeva, iSeries, etc.

Kolkata

Tel: 033-22855099, 9830048632

Fax: 033-22908152

Email: ais@duckback.co.in

Web: www.duckback.co.in

FOSS Yellow Pages

HBS System Pvt Ltd

System Integration & Service Provider Partner of IBM, DELL, HP, Sun, Microsoft, Redhat, TrendMicro, Symantec Partners of IBM for their new virtual E-commerce initiative. Solution Provider in REDHAT, ORACLE & JAVA.

New Delhi

Tel: 011-25767117, 25825001/02/03

Fax: 25861428

Email: anil@hbsinfo.com



Ingres Corporation

Ingres Corporation is a leading provider of open source database software and support services. Ingres powers customer success by reducing costs through highly innovative products that are built on open source deployment and uniquely designed for business critical applications. Ingres supports its customers with a vibrant community and world class support, globally. Based in Redwood City, California, Ingres has major development, sales, and support centers throughout the world, and more than 10,000 customers in the United States and internationally.

New Delhi

Tel: 011-40514199, Fax: +91 22 66450637

Email: sales@ingres.com; info@ingres.com

Web: www.ingres.com

Keen & Able Computers Pvt Ltd

Open Source Solutions Provider. Red Hat Ready Business Partner. Mail Servers (virt-spawn/GUI), Mailware Encryption, Clustering & Load Balancing - SAP/Oracle/MS/Thin Clients, Networks and Hub/ Monitoring, Security Consulting, Solutions, Staffing and Support.

New Delhi-110019

Tel: 011-30880046, 30880047

Mobile: 09810477448, 09891074906

Email: info@keenable.com

Web: www.keenable.com

LDS Infotech Pvt Ltd

is the authorized partner for Redhat, Linux, Microsoft, Adobe, Symantec, Oracle, IBM, Citrix etc. Software Services Offered: • Collaboration Solutions • Network Architecture • Security Solutions • Disaster Recovery • Software Licensing • Antivirus Solutions.

Mumbai

Tel: 022-25349192

Email: sales@ldsinfotech.com

Web: www.ldsinfotech.com

Pacer Automation Pvt Ltd

Pacer is leading provider of IT Infrastructure Solutions. We are partners of HP, Redhat, Cisco, VMware, Microsoft and Symantec. Our core expertise includes Consulting, building and Maintaining the Complete IT Infrastructure.

Bangalore

Tel: 060-42825001, Fax: 060-42825003

Email: solutions@pacerautomation.com

Web: www.pacerautomation.com



Red Hat India Pvt Ltd

Red Hat is the world's leading Open Source Solutions provider. Red Hat provides high-quality, affordable technology with its operating system platform, Red Hat Enterprise Linux, together with applications, management and Service Oriented Architecture (SOA) solutions, including iBee Enterprise Middleware. Red Hat also offers support, training and consulting services to its customers worldwide.

Mumbai

Tel: 022-39878880

Email: marketing-in@redhat.com

Web: www.redhat.in

Srjan Technologies Pvt Ltd

Srjan is an IT consulting company engaged in designing and building web applications, and IT infrastructure systems using open source software.

New Delhi

Tel: 011-26225006, Fax: 011-41606543

Email: business@srjan.in

Web: www.srjan.in

TechnoInfotech

A company focussed on Enterprise Solution using open source software. Key Solutions:

- Enterprise Email Solution
- Internet Security and Access Control
- Managed Services for Email Infrastructure

Mumbai

Tel: 022-06336000, Extn. 324

Email: sales@technoinfotech.com

Web: www.technoinfotech.com

Tetra Information Services Pvt Ltd

One of the leading open source providers. Our cost effective business ready solutions consist of all kind of industry verticals.

New Delhi

Tel: 011-46513113, Fax: 011-41620171

Email: sales@tetra.in

Web: www.tetra.in

Tux Technologies

Tux Technologies provides consulting and solutions based on Linux and Open Source software. Focus areas include migration, mail servers, virus and spam filtering, clustering, firewalls, proxy servers, VPNs, server optimization.

New Delhi

Tel: 011-27348104, Mobile: 09212088104

Email: info@tuxtechnologies.co.in

Web: www.tuxtechnologies.co.in

FOSS Yellow Pages

Want to register your organisation in
FOSS Yellow Pages For **FREE**

Call: Dhiraj (Delhi) 09811206582 Somalah (Bangalore) 09986075717
or mail: dhiraj.khare@efyindia.com, somalah.km@efyindia.com

*Offer for limited period.

Qt 4.5, Qt Creator and Qt SDK

Nokia has announced the availability of version 4.5 of the Qt cross-platform application and UI framework. The company also introduced Qt Creator—a new lightweight cross-platform IDE. Qt 4.5 and Qt Creator combined, comprise the Qt SDK.

Qt 4.5 includes several new features. Significant performance enhancements were made to the graphics system, data handling, and the Web engine. These improvements resulted in an appreciable performance increase in Qt-based applications. Qt 4.5 also improves integration with the WebKit rendering engine that blends Web and native content into a richer user experience. This includes Netscape Plugin API support allowing Qt applications to load Flash (such as the YouTube player); more advanced Web UI effects including animations, transformations and zooming; and new JavaScript engine. Visit www.qtsoftware.com/downloads to grab the latest version.

Wind River Linux 3.0

Wind River Linux 3.0 has been released based on Linux kernel 2.6.27 and GCC 4.3. The new version comes with more than 500 software packages, including packages for multimedia, graphics and HMI (human-machine interface) technologies, such as X.org, GTK+, GNOME or GStreamer. The new release also offers sophisticated multi-core features such as virtualisation based on KVM and multi-core offload, allowing customers to utilise the potential of modern multi-core hardware.

Novell releases SUSE Linux Enterprise 11

Novell has announced the availability of SUSE Linux Enterprise 11 (SLE 11). The new release contains major enhancements to SLE Server and SLE Desktop, and delivers two new extensions: SUSE Linux Enterprise Mono Extension, which enables customers to run fully supported .NET-based applications on Linux, and SUSE Linux Enterprise High Availability Extension, a clustering product that ensures uptime for mission-critical applications while slashing the cost of ownership for high availability systems. As part of a complete solution, SUSE Linux Enterprise 11 will be supported by Novell's global strategic partners—including Cisco, Dell, Fujitsu-Siemens, HP, IBM, Intel, Microsoft, NEC, Oracle and SAP—and its entire ecosystem of solutions, integration, hardware and software partners.

SUSE Linux Enterprise 11 runs on the leading hardware platforms, including x86-32, x86-64, Itanium, IBM POWER and IBM System z. In addition, it has been optimised to run at near-native performance on all major hypervisors, including Xen, VMware ESX and Microsoft Hyper-V.

SLE 11 includes an all-round update of its package stack; some highlights include: kernel 2.6.27, GCC 4.3.2, glibc 2.9, Samba 3.2.4, X.org 7.4, GNOME 2.24, KDE 4.1, Perl 5.10, PHP 5.2.6, Python 2.6.0, Ruby 1.8.7, Apache 2.2.10, Bind 9.5.0P2, etc. The default file system in new installations was changed from ReiserFS to ext3, while LVM2 has replaced EVMS2. For the desktop users, SLE Desktop 11 features the ability to play MP3, AAC and Windows multimedia file formats, and the latest version of OpenOffice.org Novell Edition supports a range of Microsoft Office file formats.

Evaluation editions for the products are available from download.novell.com.

AMP5071 geared for communications system engineers

Performance Technologies, a provider of innovative network communications solutions, has announced the availability of its newest IPnexus MicroTCA systems, the AMP5071. It is targeted at systems engineers developing IP-based communication products, and comes loaded with the company's NexusWare Carrier Grade Linux OS and development environment.

It comes integrated with a choice of processing Advanced Mezzanine Card (AMC). Developers can select an Intel Core 2 Duo processor, or a Freescale MPC8641D dual-core 1 GHz PowerPC processor.

The advanced features of the IPnexus AMP5071 Platform include: redundant hot-swappable power supplies and fan trays, a sturdy steel enclosure, and high quality embedded components designed on the motherboard. It also has high-performance connectivity with dual 1 Gb Ethernet links and x4 PCI Express lanes to each AMC slot and Quad 1 Gb Ethernet uplinks (rear panel access); an extensive list of optional communications protocols for immediate integration; the MicroTCA architecture to maximise space savings; and virtually endless combinations of compute, storage, connectivity options with the company's wide range of supporting AMCs. OEM quantity pricing configured with a processor AMC, storage, and NexusWare, starts at \$4,595.



FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

The gateway to the best FOSS products & services

Delhi (Delhi) 09811206582, Somalah (Bangalore) 09986075717

Veeras Infotek Private Limited

An organization providing solutions in the domains of Infrastructure Integration, Information Integrity, Business Applications and Professional Services.

Chennai

Tel: 044-42210000, Fax: 26144966

Email: info@veeras.com

Web: www.veeras.com

Linux-Based Web-Hosting

Manas Hosting

Manas Hosting is a Bangalore-based company that is dedicated in helping small and medium business companies to reach customers online. We believe that by creating a website, all you have is just web presence, but to get effective traffic on your website, it is equally important to have a well designed one. This is why we provide the best of Web Hosting and Web Designing services. Also, our services are backed with exceptionally good quality and low costs.

Bangalore

Tel: 080-42400900

Email: enquiry@manashosting.com

Web: www.manashosting.com

Linux Desktop

Openserve Infotech Pvt Ltd

Openserve Infotech Pvt Ltd (Financial Accounting & Inventory on Linux) offers a complete Linux Desktop for SME users. Its affordable (Rs. 500+) is as spacious (scheme), Friendly (Graphical User Interface) and Secure (Virus free).

New Delhi

Tel: 011-26016670-71, Fax: 26016672

Email: info@openit.com

Web: www.openit.com

Linux Experts

Intaglio Solutions

We are the leading and testing partners of Red Hat, and the first to conduct RHCE exam in Delhi for the first time ever.

New Delhi

Tel: 011-41582017, 45515765

Email: info@intaglio-solutions.com

Web: www.intaglio-solutions.com

Linux Vendor/Distributors

GT Enterprises

Authorized distributor for Red Hat and JBoss range of products. We also represent various OSS Applications and Development Tools like JBoss, VMware, Oracle DB, MySQL, CodeWeaver's JFree, Sybase, Zimbra, Zend, A PHP Company, High Performance Computing Solutions from The Portland Group, Absoft, Patheon/Logic and Intel Compilers, GnuMail Messaging solution on Linux Platform.

Bangalore

Mobile: +91-9845029030, +91-9343861758

Email: sales@gt-india.com

Web: www.gt-india.com

Taurisoft

Contact us for any Linux Distribution at reasonable rates. Members get additional discounts and Free CD/DVDs with each purchase. Visit our website for product and membership details.

Mumbai

Mobile: 09860459928, 09822697824

Email: taurisoft@gmail.com

Web: www.taurisoft.net/taurisoft.com

Software Subscriptions

Blue Chip Computers

Available Red Hat Enterprise Linux, Suse Linux Enterprise Server / Desktop, iBos, Oracle APCentis Backup, Antivirus for Linux, Veracrypt, Thanks/GeoTrust SSL Certificates and many other original software licenses.

Mumbai

Tel: 022-25001812, Mobile: 09821097258

Email: bluechip@redhat.com

Web: www.bluechip-india.com

Software Development

Carlzon Software (P) Ltd

Carlzon's flagship product is Reliance Internet Server, a complete integrated software product covering modules like mail server, proxy server, gateway, antivirus scanner, anti-spam, groupware, bandwidth aggregator & manager, firewall, chat server and tel-server infrastructure.

Chennai

Tel: 044-24650222, 0226, 0266

Email: info@carlzon.com

Web: www.carlzon.com

DeepRoot Linux Pvt Ltd

DeepRoot Linux is a customised old GNU/Linux and Free Software company based in Bangalore. We develop Free Software products that are quick to deploy and easy to use.

Bangalore

Tel: 080-40460000

Email: shivanand@deeproot.in

Web: www.deeproot.in



InfoAxeon Technologies Ltd

InfoAxeon designs, develops and supports enterprise solutions, stacks, leveraging open standards and open source technologies. InfoAxeon's focus areas are Enterprise Integration, CRM, Content & Knowledge Management and e-Learning.

Needs

Tel: 0120-4350040, Mobile: 09810425760

Email: sales@infoaxen.com

Web: http://opensource.infoaxen.com

Integra Micro Software Services (P) Ltd

Integra focuses on providing professional services for software development and IP generation to customers. Integra has a major practice in offering Telecom Services and works for Telecom companies, Device Manufacturers, Networking companies, Semiconductor and Application development companies across the globe.

Bangalore

Tel: 080-28565801/05, Fax: 080-28565800

Email: tpvarun@integratms.com

Web: www.integratms.com

Webtune.com Pvt Ltd

Webtune.com is your one-stop, total web site support organization. We provide high-quality website services and web-based software support to any kind of websites, irrespective of the domain of the industry segments.

Bangalore

Tel: 080-4115 2929

Email: santosh@webtune.com

Web: www.webtune.com

Want to register your organisation in
FOSS Yellow Pages
For **FREE**
CALL: DELHI (DELHI) 09811206582
BANGALORE (BANGALORE) 09986075717
Still in order web

Unistal Systems Pvt Ltd

Unistal is pioneer in Data Recovery Software & Services. Also Unistal is national dealer & support partner for BitDefender Antivirus products.

New Delhi

Tel: 011-26288580, Fax: 011-26210066

Email: itakes@unistal.com

Web: www.unistal.com

Software and Web Development

Beas eArchitect Integrated Services Pvt Ltd

Application Development, Web Design, SEO, Web Marketing, Web Development.

New Mumbai

Tel: 022-27821617, Mobile: 9820156561

Fax: 022-27821617

Email: info@beasarchitect.com

Web: www.beasarchitect.com

CATEGORIES FOR FOSS YELLOW PAGES

- Consultants
- Consultant Firm
- Embedded Solutions
- Enterprise Communication Solutions
- High Performance Computing
- IT Infrastructure Solutions
- Linux-based Web-hosting
- Mobile Solutions
- Software Development
- Training for Professionals
- Training for Corporate
- Third-Party Solutions

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Dhraj (Delhi) 09811206582, Somalah (Bangalore) 09986075717

Mr Site Takeaway Website Pvt Ltd

Our product is a unique concept in India using which a person without having any technical knowledge can create his website within 1 hour. We also have a Customer Care Center in India for any kind of other sales help. We are already selling it world over with over 66,000 copies sold. It comes with FREE Domain Name, Web Hosting, and Customer Care Center for Free Support via Phone and Email and features like PayPal Shopping Cart, Guestbook, Photo Gallery, Contact Form, Forums, Blogs and many more. The price of complete package is just Rs 2,999 per year.

Patna

Mobile: 91-9790531882

Email: pankaj@mr-site.co.in

Web: www.mr-site.co.in

Salah Software

We are specialized in developing custom strategic software solutions using our solid foundation on focused industry domains and technologies. Also providing superior Solution Edge to our Clients to enable them to gain a competitive edge and maximize their Return on Investments (ROI).

New Delhi

Tel: 011-41648288, 69001566

Email: ceo@salahsoftware.com

Web: www.salahsoftware.com

Thin Client Solutions

Digital Waves

The System Integration business unit offers end-to-end solutions on Desktops, Servers, Workstations, HPC Clusters, Router/Firewall, Networking, Security/Surveillance & Enterprise Storage. With our own POWERIX branded range of Products, we offer complete solutions for Animation, HPC, Clusters, Storage & Thin Client Computing.

Mobile: 09880715253

Email: range@digitalwaves.in

Web: www.digitalwaves.in

Enjoy Network Solutions

Gujarat based Thin Client Solution Provider. Providing Serial & Parallel Thin Client PCs & Full Featured Thin Client OS to perfectly suit needs of different working environment.

Active Dealer Channel all over India.

Gujarat

Tel: 0260-3203400, 3241732, 3251732

Mobile: 09077107650, 09886007650

Email: info@enjoyworld.com

Web: www.enjoyworld.com

Training for Corporate

Bascom Bridge

Bascom Bridge is Red Hat Certified partner for Enterprise Linux 5 and also providing training to the individuals and corporate on open source technologies like PHP, MySQL etc.

Ahmedabad

Tel: 079-27545455-46

Fax: 079-27545486

Email: info@bascombridge.com

Web: www.bascombridge.com

Brainnet

Kolkata

Tel: 033-40076450

Email: brainnet@brainware-india.com

Web: www.brainware-india.com

Centre for Excellence in Telecom Technology and Management (CETTM), MTRIL

MTRIL's Centre for Excellence in Telecom Technology and Management (CETTM) is a state of the art facility to impart Technical, Managerial and corporate training to Telecom Management personnel. CETTM has AC lecture halls, computer Lab and residential facility.

Mumbai

Tel: 022-25714500, 25714566, 25714565, 25714566

Fax: 022-25706700

Email: contact@cettm.mtril.in

Web: <http://cettm.mtril.in/index>

Complete Open Source Solutions

Hyderabad

Tel: 040-66773365, 9849742065

Email: sayak.sajjan@gmail.com

Web: www.cossindia.com

ElectroMech

Redhat Linux and open source solution, RHCE, RHCSA training and exam center/Ahmedabad and vadodra.

Ahmedabad

Tel: 079-40027998

Email: electromech@electromech.info

Web: www.electromech.info



FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

Focus Infotech

Focus Infotech Advanced Education is the quality symbol of high-end Advanced Technology Education in the state. We are providing excellent services on Linux Technology, Training, Certifications and live projects to students and corporates, since 2000.

Cochin

Tel: 0484-2336324

Email: enquiry@focusinfotech.com

Web: www.focusinfotech.com

Gujarat Infotech Ltd

GIL is a IT company with 17 years of experience in computer training field. We have experienced and certified faculty for the open Source courses like Redhat, Ubuntu, and PHP, MySQL.

Ahmedabad

Tel: 079-27452276, Fax: 27414250

Email: info@gujaratinfotech.com

Web: www.gujaratinfotech.com

Lynux Academy Pvt Ltd

India's premier Linux and OSE training Institute.

Chennai

Tel: 044-42171278, 9849886661

Email: contact@lynuxacademy.com

Web: www.lynuxacademy.com

Linux Learning Centre Private Limited

Pioneers in training on Linux technologies.

Bangalore

Tel: 080-22428538, 20000838

Email: info@linuxlearningcentre.com

Web: www.linuxlearningcentre.com

Maze Net Solutions (P) Ltd

Maze Net Solution (P) Ltd, is a pioneer in providing solutions through on time, quality delivered services in the fields of ERP, Software and Networking.

while providing outstanding training to aspiring IT Professionals and Call Center Executives. Backed by a team of professional workforce and global alliances, our prime objective is to offer the best blend of technologies in the sphere of Information Technology (IT) and Information Technology Enabled Services (ITES).

Chennai

Tel: 044-45582525

Email: info@maizesolutions.com

Web: www.maizesolutions.com

Netweb Technologies

Simplified and scalable storage solutions.

Bangalore

Tel: 080-41142665, 32719616

Email: info@netwebindia.com

Web: www.netwebindia.com

New Horizons India Ltd

New Horizons India Ltd, a joint venture of New Horizons Worldwide, Inc. (NASDAQ: NHHM) and the Shriram group, is an Indian company operational since 2002 with a global footprint engaged in the business of knowledge delivery through acquiring, creating, developing, managing, lending and licensing knowledge to the areas of IT, Applied Learning, Technology Services and Supplementary Education. The company has pan India presence with 15 offices and employs 760 people.

New Delhi

Tel: 011-43612400

Email: info@nhindia.com

Web: www.nhindia.com

FOSS Yellow Pages

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To: yellowpages@openTIS.com, Please! Contact:

Dhraj (Delhi) 09811206582, Somalah (Bangalore) 09986075717

Network NUTS

India's only Networking Institute by Corporate Trainers. Providing Corporate and Open Access for RHCE / RHCPS training and certification. Conducted 250+ Red Hat exams with 95% pass in last 3 months. The BEST in APAC.

New Delhi

Tel: 46526980-2

Mobile: 09310024903, 09312411592

Email: info@networknuts.net

Web: www.networknuts.net

STG International Ltd

An IT Training and Solution Company. Over an experience of 14 years. We are ISO 9001:2000 Certified Authorized Training Partners of Red Hat & IBM-CEE. We cover all Software Trainings.

New Delhi

Tel: 011-40560941-42, Mobile:

09873108801

Email: rahil@stgin

Web: www.stgindia.com

www.stgglobal.com

TMS Institute of Information Technology Pvt Ltd

Join Red Hat training and get 100% job guarantee. World's most respected Linux certification. Also Red Hat training, you are back to join as a Linux Administrator/Offshore Engineer.

New Delhi

Tel: 011-3085100, Fax: 30851103

Email: nks@tms.co.in

Web: www.tms.co.in

Webel Informatics Ltd

Webel Informatics Ltd (PIL), a Government of West Bengal Undertaking. We do Red Hat Training Partner and CCSC-Regulated Networking Academy. We conduct RHCE, RHCSA, CCNA, Hardware and Software course.

Kolkata

Tel: 033-22833600, Mobile: 09433111110

Email: enquiry@webelinformatics.com

Web: www.webelinformatics.com

Training for Professionals

AEM

AEM is the Best Certified Red Hat Training Partner in Eastern India since last 3 years. AEM conducted more than 500 RHCE exams with 95-100% pass rate. Other courses—

RHCE, CCNA, MCSA, CCNA

Kolkata

Tel: 033-25488736, Mobile: 09899075018

Email: webelinfo1@gmail.com

Web: www.aemk.org

Agam Institute of Technology

In Agam Institute of Technology, we provide hardware and networking training since last 10 years. We specialise in open source operating systems like Red Hat Linux since we are their preferred training partners.

Dehradun

Tel: 0135-2673712, Mobile: 09793099660

Web: www.agamtechindia.com

Amritia Institute of Computer Technology

Amritia Technologies provides an enterprise training & high end certification programs and networking solutions like Red Hat Linux, Red Hat Security Services, Cisco, Sun Solaris, Cyber Security Program IBM AIX and we are with a strong focus on quality standards and proven technology processes with most profound principle of Love and Softest Service.

Mobile: 098083733174

Email: sict.hydr@amritia.ac.in

Web: www.amritiahyd.org

Centre For Industrial Research and Staff Performance

A Unique Institute catering to the need for industries as well as Students for trainings on IT, CCSC certification, PLC, VLSI, ACAD, Pneumatics, Behavior Science and Handcraft.

Bhopal

Tel: 0755-2661412, 2661550

Fax: 0755-4220002

Email: crisp@crispindia.com

Web: www.crispindia.com

Center for Open Source Development And Research

Linux, Open source & embedded system training institute and development. All trainings provided by experienced experts & administrators only. Quality training (corporate and individual). We expertise in open source solution. Our cost effective, business ready solutions caters of all kind of industry verticals.

New Delhi

Mobile: 09512508496

Email: info@cfosdr.com

Web: www.cfosdr.com

Disconet Infotech (P) Ltd

Authorized Red Hat Study cum Exam Centre. Courses Offered: RHCE, RHCSA, CCNA, MCSA.

Kolkata

Tel: 033-25365608, Mobile: 09831765913

Email: info@disconetinfo.com

Web: www.disconetinfo.com

CMS Computer Institute

Red Hat Training partner with 3 Red Hat Certified Faculties, Cisco Certified (CCNP) Faculty, 3 Microsoft Certified Faculties having state of the art IT Infrastructure Flexible Batch Timings Available. Leading Networking Institute in Madhya Pradesh.

Bhopal

Tel: 0540-3250604, 6821775

Email: bhopal@cmscomputerinstitute.co.in

Web: www.cmscomputerinstitute.co.in

Cyber Max Technologies

QSS Solution Provider Red Hat Training Partners, Oracle/IBM Thin Clients, Networking and Security Consultancy. Also available CCNA and Oracle Training on Linux. Also available Laptops & PCs.

Bikaner

Tel: 0151-2262105, Mobile: 09923173260

Email: cmtech.bikaner@gmail.com

kgupta.ashish@gmail.com

Disha Institute

A franchisee of Unisoft Technologies. Providing IT Training & Computer Hardware & Networking.

Dehradun

Tel: 3208054, 09867168902

Email: thedishainstitute@gmail.com

Web: www.unisofttechnologies.com

EON Infotech Limited (TECHNOSCHOOL)

Technoschool is the most Reputable Training Centre for Red Hat Linux (Open Source) in the Northern Region. We are fully aware of the industry requirements for Consultants are from Linux industry. We are committed to make you a total industry ready individual so that your dreams of a professional career are fulfilled.

Chandigarh

Tel: 0172-5067665-67, 2006049

Fax: 0172-2615465

Email: info@technoschool.net

Web: <http://technoschool.net>

GT Computer Hardware Engineering College (P) Ltd

Imparting training on Computer

Hardware Networking, Mobile

Phone Maintenance & International Certifications

Japur

Tel: 0141-3213378

Email: harddisa_gt@education.net

Web: www.gteducation.net

HCL Career Development Centre Bhopal

As the forerunner of the most significant pursuit of human mind (IT), HCL strongly believes, "Only a leader can transform you into a Leader". HCL CDC is a formalization of the experience and wisdom which has been perfected over three decades.

Bhopal

Tel: 0755-404852

Email: bhopal@hclcdc.in

Web: www.hclcdc.in

INZTRIX E Technologies Pvt Ltd

No. 1 Training provider in the region nearest

Tel: 0121-4020111, 4020222

Mobile: 09627966664

Email: info@inztrix.com

Web: www.inztrix.com

Indian Institute of Job Oriented Training Centre

Amrohtabad

Tel: 079-49072244—2255—2266

Mobile: 09806789595

Email: info@iijt.net

Web: www.iijt.net

Institute of Advance Network Technology (IANT)

• Hardware Engg • Networking • Software Engg • Multimedia Training

Amrohtabad

Tel: 079-32516577, 29607730

Fax: 079-26607730

Email: contact@iantindia.com

Web: www.iantindia.com

IPOC

Bridging Gap with professionals

Lucknow

Tel: 0522-3919496

Email: ipoclw@yahoo.co.in

Web: www.ipoc.co.in

Koenig Solutions (P) Ltd

A reputed training provider in India. Authorized training partner of Red Hat, Novell and Linux Professional Institute. Offering training for RHCE, RHCSA, CUP, CUE, LPI 1 & 3.

New Delhi

FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To: chiraj.khare@openTls.com or somaiah.km@etyindia.com

Chiraj (Delhi) 09811206582, Somaiah (Bangalore) 09986075717

Mobile: 09810710143, **Fax:** 011-25886000
Email: info@knoxig-solutions.com
Web: www.knoxig-solutions.com

NACS/CIT

We are Providing Training of LINUX to Professionals & Engineers

Mumbai

Tel: 0121-2420687, **Mobile:** 9897526688
Email: info@nacsglobal.com
Web: www.nacsglobal.com

NACS Infosystems (P) Ltd

NACS is a organization which is providing training for all international certification, and also NACS is the authorized Training Partner of Redhat and also having learning centre of THOMSON PROMETRIE and PEARSON VUE

Mumbai

Tel: 0121-2767796, **Fax:** 0121-4006651
Mobile: 09897799003
Email: info@nacsglobal.com, naish@nacsglobal.com
Web: www.nacsglobal.com

Netdix Computing Systems


We are one-of-a-kind center for professional and freetime school focusing on ground breaking technology development ground distributed systems, network, storage networks, virtualization and fundamental algorithms optimized for various appliances

Bangalore

Tel: 080-29640700
Mobile: 09740849885
Email: info@netdix.co

NetMax Technologies

Training Partner of Redhat, Cisco
Chandigarh
Tel: 0172-2506351, 3216555
Email: mail.netmax@gmail.com
Web: www.netmaxtech.com



Netxpri Institute of Advanced Networking

Netxpri India is a leading organization to provide Open Source training on Redhat Linux RHCT and RHCE Training with 30 hrs extra exam preparation module

Noida

Tel: 0120-4346547, **Mobile:** 00268229612
Email: anup.vyas@netxpriindia.com
Web: www.netxpriindia.com

Netzone Intektech Services Pvt Ltd

Specialty switches for MDSB, CDNA and RHCE on RHEL 5 with exam prep module on fully equipped lab including IBM servers, 20+ routers and switches etc. Weekend batches are also available

New Delhi

Tel: 011-46015674, **Mobile:** 9212114211
Email: info@netzoneindia.net

Neuron IT Solutions

We offer end to end services and support to implement and manage your IT infrastructure needs. We also offer Consulting services and Training in Advanced Linux Administration

Chennai

Mobile: 09700054948
Email: enquiry@neuronit.in
Web: www.neuronit.in

Plexus Software Security Systems Pvt Ltd

Plexus, incorporated in January 2008 is successfully emerged as one of the best IT Company for Networking, Messaging & Security Solutions and Security Training. Networking, Messaging & Security solutions is coupled with the expertise of its training, the has put Plexus in the unique position of deriving synergies between Networking, Messaging & Security Solutions and IT Training

Chennai

Tel: 044-2433 7355
Email: training@plexus.co.in
Web: www.plexus.co.in

Professional Group of Education

RHCE & RHCSA Certifications
Jabalpur
Tel: 0761-4030376
Mobile: 09425152831
Email: nishi.vikas@gmail.com

Q-SOFT Systems & Solutions Pvt Ltd

Q-SOFT is in a unique position for providing technical training required to become a Linux Administration professional. Since inception, the commitment of Q-SOFT towards training is outstanding. We Train in Sun Solaris, Suse Linux & Redhat Linux

Bangalore

Tel: 080-26630207, 26544135, 22440507
Mobile: +91 9345 22834
Email: customers@qsoftindia.com
Web: www.qsoftindia.com

Software Technology Network

STN is one of the most acknowledged names in Software Development and Training. Apart from providing Software Solutions to various companies, STN is also involved in imparting High-end project based training to students of MCA and B.Tech etc. of various institutes.

Chandigarh

Tel: 0172-5086629
Email: stn2001@rediffmail.com
Web: stntechnologies.com

South Delhi Computer Centre

SDCC is not providing technical training courses (software, hardware, networking, graphics) with career courses like OCEANO 101 and 102 Level and B SCIT (M SCIT)M SCIT (M SCIT) from KARNATAKA STATE OPEN UNIVERSITY

New Delhi

Tel: 011-26183327, **Fax:** 011-26143642
Email: southdelhicomputercentre@gmail.com, southdelhicomputercentre@hotmail.com
Web: www.itshiksha.com
www.itshiksha.org

Systems Quest

Making IT professionals TODAY

Bangalore

Tel: 080-41301814
Email: directors@systemsquest.com
Web: www.systemsquest.com

Trinax FuturePerfect

A Div of Trinax IT Infrastructure and Services Limited. Redhat RHCE, RHCT Training & Exam Center, MCIT, MCITP, MCSB CB, CDNA, CDNF, Pinnacle, Certix

Mumbai

Tel: 022-40681513, **Mobile:** 0987705638
Fax: 022-40681001
Email: futureperfect@trinax.in
Web: www.trinax.in

Vibrant e Technologies Ltd

Vibrant e Technologies Ltd. is a authorized Red Hat Test and Training Centre, has won the prestigious award "REDHAT BEST CERTIFIED TRAINING PARTNER 2007-2008" for Western region. Vibrant offers courses for RHCE & RHCSA etc.

Mumbai

Tel: 022-26285096/6791
Email: etbnet@vsnl.net
Web: www.vibrantelearning.com

Ultramax Infonet Technologies Pvt Ltd

Training in IT related courses, authorized testing center of Eminent, AUE and Red Hat

Mumbai

Tel: 022-47660217
Email: umesh.raut@ultramaxit.com
Web: www.ultramaxit.com

Yash Infotech

Authorized Training & Exam Center
Best Performing Center in Lucknow for RH Training and Examinations
Linux & Open Source training
Institute for IT professionals & Corporate Offering Quality Training for RHCE, RHCSA, PHP, Shell Script, Virtualization and Troubleshooting Techniques & Tools

Lucknow

Tel: 0522-4043386, **Fax:** 0522-4043386
Email: yashinfotech.lko@gmail.com

Want to
register
your
organisation
in FOSS
Yellow
Pages For

FREE

Call: Chiraj (Delhi) 09811206582

Somaiah (Bangalore) 09986075717

or mail: chiraj.khare@etyindia.com

or somaiah.km@etyindia.com

Circle 10 on Reader Service



SUBSCRIBE TODAY & SAVE upto 40%



Magazine	Unit Price (Rs)	5 Years (Rs)	3 Years (Rs)	1 Year (Rs)
Electronics For You (EFY)	40/-	1445 <input type="checkbox"/>	1010 <input type="checkbox"/>	385 <input type="checkbox"/>
Electronics For You (with CD)	60/-	2160 <input type="checkbox"/>	1510 <input type="checkbox"/>	575 <input type="checkbox"/>
Electronics Bazaar	50/-	— <input type="checkbox"/>	1260 <input type="checkbox"/>	480 <input type="checkbox"/>
Linux For You (with CD & DVD)	100/-	3600 <input type="checkbox"/>	2520 <input type="checkbox"/>	960 <input type="checkbox"/>
BenefIT	50/-	— <input type="checkbox"/>	1250 <input type="checkbox"/>	475 <input type="checkbox"/>
Facts For You	100/-	3590 <input type="checkbox"/>	2510 <input type="checkbox"/>	950 <input type="checkbox"/>

- ✓ Don't Miss A Single Issue!
- ✓ Ensure Regular Supply
- ✓ FREE Replacement Policy*
- ✓ Get Your Money Back if Not Satisfied**
- ✓ Order Now For FREE Home Delivery

* Replacement will be made upon receipt of the damaged / non-receipt copies within 30 days of its publication.
** After three months, if you are not satisfied with the magazine, your balance amount will be returned.

Name _____ Company Name _____ Designation _____

Mailing Address _____ Distt _____

City _____ Pin Code _____ State _____ Phone _____

Email _____ Subscription No. (for existing subscribers only) _____

I would like to subscribe to the above (✓) marked magazine(s), for the period _____ Please find enclosed a sum of Rs _____ by DD/MO/cashed cheque* bearing the No. _____ dt. _____ in favour of EFY Enterprises Pvt Ltd, payable at Delhi. (*Please add Rs 50 as non-metro cheque) or charge my credit card ☐ VISA ☐ MasterCard Please charge Rs. _____ against my credit card No. _____

Date of Birth ____/____/____ (dd/mm/yy) Card Expiry Date ____/____/____ (mm/yyyy) Signature (as on the card) _____

Terms:- # These rates are applicable for new subscribers as well as renewal by existing subscribers. # The rates are valid for subscribers within India only. # Please allow 4-6 weeks for processing of your subscription. # Disputes, if any, are subject to exclusive jurisdiction of competent courts and forums in Delhi/New Delhi only.

Send this filled-in form or its photocopy to :



EFY Enterprises Pvt Ltd D-87/1, Okhla Industrial Area, Phase 1, New Delhi 110 020; Ph: 011-26810601-03; Fax: 011-26817563; e-mail: info@efyindia.com; website: www.efyindia.com

Microsoft

BE LOOKED UP TO

Be the one to save the day for developers with a solution, or share something new to become the one everyone would talk about. Start a new project, join an existing one, or download softwares created by fellow developers.

ENTER A NEW HALL OF FAME AT WWW.CODEPLEX.COM,
THE MICROSOFT OPEN SOURCE PROJECT HOSTING WEBSITE.

CodePlex



McGraw-Hill

Technology News

Got Anjal in your netbook?

Hey, you'd love to have it on your PC too! Srinivasa Ragavan V writes in his blog [blogs.gnome.org/sragavan/2009/03/18/announcing-anjal-the-new-mail-for-netbooks]. "Anjal is a new mail UI created on top of Evolution. It has a very interesting UI and features that would make it very suitable for low memory/processor/resolution devices. It is Evolution and EDS at the back of Anjal." He says that the Evolution mail library has been broken down to a smaller component to make a reusable, shared library. Besides this, Anjal, which means 'mail' in Tamil, will use Webkit for mail rendering and composing.

He writes, "It features a nice multi-line message list, with unquoted text-preview of the latest messages in the thread. It's configurable to make it even lighter, by choosing an sqlite-cursor-based tree view, which would have the visible window of the message list on memory, which is very suitable for mobile environments. It supports tabbed browsing, a small preference window and new account set-up." There is a very interesting feature, Conversation View, in which you can see your own e-mail as part of the thread. This is something that many people are used to as Gmail users, and take pains to configure in their desktop mail client application, which is seldom achievable in most currently popular clients.

The project is only a month old and you will notice "...more changes as it grows, w.r.t. features and UI, to make it a good mail-client for GNOME on Netbooks." The code is hosted at git://git.gnome.org/anjal, and the developers are hard at work to release version 0.1 pretty soon.

The new PCLinuxOS after more than a year

The PCLinuxOS 2009.1 has been released featuring kernel 2.6.26.8, KDE 3.5.10, OpenOffice.org 3.0, Firefox 3.0.7, Thunderbird 2.0.0.14, KTorrent, Frostwire, Amarok, Flash, Java JRE, Compiz-Fusion 3D, etc. KDE4 will be offered as an alternative desktop environment from the repository, as soon as the developers are able to stabilise it. Separately, the pclos-gnome team has announced the 2009.1 release of PCLOS-GNOME featuring GNOME 2.24.3, GNOME Office

Apps, Brasero, and Exaile Music Player, besides the regular tools.

Both releases contain almost 2.5 GB of software compressed on a bootable live CD that can be installed to your hard drive, provided it is compatible with your system and you like the distribution. Over 8000 additional packages are available after the hard drive install through the Synaptic Software Manager. Visit www.pclinuxos.com for more details.



A complete portfolio of JBoss Enterprise runtimes and tools

Red Hat has released JBoss Developer Studio 2.0-Portfolio Edition. According to Red Hat, this complete development environment not only enhances the Eclipse-based toolset, but also adds JBoss Enterprise Application, Portal, SOA and Data Services Platforms. The solution provides a robust, integrated development environment for rich Web applications, mission-critical enterprise applications, integration services in a single package.

The JBoss Developer Studio builds on Eclipse, the de-facto standard Java developer toolset, and includes a comprehensive set of tools that developers need to quickly create highly interactive applications and services using technologies like Seam, Java EE, AJAX and Portlets as part of standards-based SOA. The ability to quickly deploy enterprise applications is important for organisations looking to achieve value through cost and time savings.

JBoss Developer Studio Portfolio Edition is supposed to help reduce the time between development and deployment by including and supporting the full portfolio of JBoss Enterprise runtimes and tools. These include JBoss Developer Studio 2.0, JBoss Enterprise Application Platform, JBoss SOA Platform, JBoss Enterprise Data Services Platform, JBoss Enterprise Portal Platform, JBoss Operations Network, OpenJDK, and RHEL. For more information on JBoss Developer Studio - Portfolio Edition, please visit www.jboss.com/products/devstudio.



National Conference on Open Source Software

May 25th - 26th, 2009, Mumbai, India

Organised by
C-DAC, Mumbai

Supported by
IEEE Computer Society, Mumbai and Chennai chapters
Computer Society of India, Div II on Software & SIG-OSS

Call For Papers

NCOSS—2009 is a forum to bring together the various groups working on developing Open Source Applications catering to specific domains in the ICT world—education, health, accessibility, localisation, e-commerce, disaster management, expert systems, machine learning, etc. A number of high-quality software solutions are available in many of these areas, for example, SugarCRM, Koha, Drupal, Moodle, Sahana, CollabCAD, etc. Work on these systems require a combination of domain knowledge and development expertise. Much of the public awareness in open source is focussed on desktop, operating system and general productivity tools. With this background, NCOSS-09 has chosen to focus on the layer above this, bringing together groups working on various application domains.

The conference will present experiences in deploying FOSS applications, comparative studies among competing software solutions, efforts in adapting and localising FOSS applications, development of new applications, etc. The conference will consist of the following:

- Invited talks by experts from India and abroad
- Presentation of contributed papers selected based on refereeing by a panel of referees
- Exhibition by industry and academia
- Pre conference Tutorials (on May 24th)
- Panel discussion

TOPICS

Papers are invited on the topics listed below: (Other application areas may also be considered).

Accessibility	Machine Learning and Data Mining	e-Governance
Indian Language Computing	e-Health	e-Commerce
Localisation	Knowledge Management	Disaster Management
e-Learning	Collaboration Technologies	Content Management
Information Extraction and Retrieval		

INSTRUCTIONS

- Papers must report original work carried out by the authors. The work can include enhancing existing Open Source applications for specific requirements, development of new solutions and comparative analysis of competing solutions. Direct survey or overview papers are not acceptable.
- Length should not exceed 10 pages of A4 size in length (approx. 5000 words) including figures, etc.
- Papers should be in English.
- An abstract of about 100-200 words and the area(s) under which the paper can be categorized, must be included with the paper.
- The author names and affiliations along with the main area of the paper should be given only on a separate cover sheet. Papers should be in one of the following formats: PDF, RTF or ODT. Accepted papers will be published in the conference proceedings.

Visit <http://ncoss.cdacmumbai.in> for more details and paper submissions.

Media Partner



Dr Wietse Venema and CC win annual free software awards

The Free Software Foundation has announced the winners of the annual free software awards. Creative Commons was honoured with the Award for Projects of Social Benefit, and Dutch programmer and physicist, Dr Wietse Venema, was honoured with the Award for the Advancement of Free Software. FSF founder and president Richard Stallman presented the awards.



The FSF Award for Projects of Social Benefit is presented annually to a project that intentionally and significantly benefits society by applying free software, or the ideas of the free software movement.

Since its launch in 2001, Creative Commons has worked to foster a growing body of creative, educational and scientific works that can be shared and built upon by others. Creative Commons has also worked to raise awareness of the harm inflicted by increasingly restrictive copyright regimes. In receiving this award, Creative Commons joins previous winners Groklaw (2007), Sahana (2006), and Wikipedia (2005).

The Award for the Advancement of Free Software went to Wietse Venema. The awards committee honoured both Venema's significant and wide-ranging technical contributions to network security, and his creation of the Postfix e-mail server. He joins a distinguished list of previous free software award winners: Harald Welte (2007), Ted Ts'o (2006), Andrew Tridgell (2005), Theo de Raadt (2004), Alan Cox (2003), Lawrence Lessig (2002), Guido van Rossum (2001), Brian Paul (2000), Miguel de Icaza (1999), and Larry Wall (1998).

This year's award committee was composed of Suresh Ramasubramanian (chair), Peter H. Sahas, Raj Mathur, Hong Feng, Andrew Tridgell, Jonas Oberg, Verner Vinge, Richard Stallman, and Fernanda G. Weiden.

SpikeSource targets ISVs in India

SpikeSource, along with its distribution partners, has announced plans to target independent software vendors (ISVs) in India with the 'SpikeSource Solutions Factory' platform. This is an end-to-end automated platform for testing, packaging, certifying and supporting software solutions of all kinds (including open source, proprietary and hybrid solutions) for global distribution.

"In India, ISVs are developing innovative solutions that have significant market potential, yet most of them are achieving only a fraction of their total market potential. The key obstacles are lack of access to global distributions, limited market awareness, customer reluctance to purchase solutions from unknown providers, and customer uncertainty about the quality of the solutions. The SpikeSource Solutions Factory platform helps remove these obstacles," said Ramesh Shastri, managing director, India, SpikeSource. "Along with our distribution partners, the SpikeSource Solutions Factory platform enables ISVs to deliver higher quality, more trustworthy solutions, reduce their development and support costs, and increase their sales and profit potential through access to global distributions."



MontaVista adopts Moblin

MontaVista Software has announced support for the Moblin software platform. By supporting Moblin and Intel Atom processors, developers of embedded devices can rely on the commercial quality and support provided by MontaVista Software to deliver new and innovative products to market, sources at the company said.

"We are excited to be part of the Moblin community. It offers compatibility with existing and future x86 development environments and tools that allow the reuse of existing, mature development infrastructure. By delivering MontaVista's embedded Linux commercialisation based on Moblin, device manufacturers can bring commercial devices to market quickly and cost effectively," said Joerg Bertholdt, VP of marketing at MontaVista Software.

"Since Moblin is optimised to deliver on what embedded customers need, MontaVista is well positioned to provide the time to market advantage for a wide range of customers across embedded market segments," said Ram Peddibhotla, open source business director, Intel Software and Services Group.

Wind River, Intel to market optimised multi-core solutions

Wind River has announced that working in conjunction with Intel, the two companies will collaborate to market optimised multi-core solutions for the embedded market by aligning research and development, sales and marketing, professional services, and engineering resources. The two companies will initially target the aerospace and defence, network infrastructure, industrial, medical and print imaging market segments.

More information about Wind River's multi-core solutions can be found at www.windriver.com/multicore-software.

Open Source firms selected for new trading desk solution

dbConcert has introduced a complete trading desk solution that gives the financial services industry a way to automate a majority of business functions done by most trading desks today. The solution, based entirely on open source products from Ingres, Marketcetera, and Pentaho, allows organisations to manage their business effectively and make better investment decisions—all at a fraction of the cost of proprietary software offerings.

"The days of the five million dollar projects are over," said Stephen Ferrando, chief executive officer, dbConcert. "No CIO has space in his budget to write that cheque. People are looking to open source to continue to innovate, to continue with new projects, without writing a bigger cheque. No one wants to spend if they don't know it will work. Open source gives you the ability to try it out without locking yourself in to millions of dollars of payments in the coming years."

Marketcetera, an open source trading system provider, recently launched the first open source trading platform. dbConcert took the Marketcetera solution and combined business intelligence (BI) from Pentaho, and the database from Ingres to create a complete open source solution framework built on Eclipse. The offering enables end users to execute orders electronically, track order flow, analyse market positions, and measure profit and loss.

Intel notifies AMD of cross-licence breach

Intel claims to have notified Advanced Micro Devices (AMD) that it believes AMD has breached a 2001 patent cross-licence agreement with Intel. Intel believes that Global Foundries is not a subsidiary under terms of the agreement and is therefore not licensed under the 2001 patent cross-licence agreement. Intel also said the structure of the deal between AMD and ATIC breaches a confidential portion of that agreement. Intel has asked AMD to make the relevant portion of the agreement public, but so far AMD has declined to do so, state sources in Intel. AMD's breach could result in the loss of licences and rights granted to AMD by Intel under the agreement.

"Intellectual property is a cornerstone of Intel's technology leadership and for more than 30 years, the company has believed in the strategic importance of licensing intellectual property in exchange for fair value. However, AMD cannot unilaterally extend Intel's licensing rights to a third party without Intel's consent," said Bruce Sewell, senior vice president and general counsel for Intel. "We have attempted to address our concerns with AMD without success since October. We are willing to find a resolution but at the same time we have an obligation to our stockholders to protect the billions of dollars we've invested in intellectual property," he added.

Under terms of the licence agreement, the notification to AMD means the parties will attempt to resolve the dispute through mediation. In response to the notification, AMD claimed Intel breached the agreement by notifying AMD of its breach. Intel believes that this position is inconsistent with the dispute resolution process outlined in the original agreement.

For the want of an OSS in-vehicle infotainment platform

Leading automobile manufacturers and suppliers announced the formation of the GENIVI Alliance, a non-profit organisation committed to driving the development and broad adoption of an open source In-Vehicle Infotainment (IVI) reference platform. The new alliance will unite industry-leading automotive, consumer electronics, communications and application development companies investing in the IVI market. The effort will result both in reducing time-to-market and total cost of ownership.

GENIVI Alliance founding members—the BMW Group, Delphi, General Motors Corp., Intel, Magneti Marelli, PSA Peugeot Citroën, Visteon Corp., and Wind River—are collaborating to create a shared GENIVI platform, a common software architecture that is scalable across product lines and generations. The GENIVI platform will accelerate the pace at which auto-makers can deliver new solutions, bringing them closer to the lifecycle of consumer devices, and accelerating new business models, such as connected services.

"GENIVI will challenge the traditional approach of proprietary solutions and spawn a level of creativity not yet seen in this segment," said Graham Smethurst, GENIVI spokesperson and BMW group general manager, Infotainment and Communication Systems. "Collaborating on a common reference platform in non-differentiating areas of the architecture will allow GENIVI members to focus on the development and integration of innovative functionality."





KNOW HOW



Q How can I get the latest installation Live CDs of Xubuntu, Kubuntu and Ubuntu.
—Noel J Purtado, by e-mail

These distros are free and can be downloaded from www.ubuntu.com/getubuntu/download. The installer CD is around 700 MB. If you do not have a good Internet connection, then you can send a request for a free CD at shipit.ubuntu.com. This service generally takes 6-10 weeks to deliver your CD.

Q I have downloaded some files from the Internet and they are in the .uif format. Please let me know how to open these files. I am using Fedora on my desktop.
—Tribhuvan, by e-mail

The file format that you have mentioned is created by the Magic ISO, a Windows application that can create, edit, and extract CD/DVD image files. For Windows, there is a freeware tool called MagicDisk that can open UIF images, but in Linux you need to convert it to ISO image files to open it. The fastest way to convert a UIF image to an ISO image is to use UIF2ISO. You can download

and install this command line tool from aluisgi.altervista.org/mytools/uif2iso.zip.

After downloading, unzip it and navigate into the *src* directory.

```
unzip uif2iso.zip
cd src
```

Now compile the source by following the commands shown below:

```
make
make install
```

Once you are over with the installation of the program, follow the steps to open your UIF image.

- To open the file, first convert the .uif file to an .iso with the following command:

```
uif2iso uif2iso.exe myuif.uif
```

- Now burn the ISO image with K3b or any other CD burning tool, or mount the ISO image with the following command:

```
mount -o loop  
-o ro,loop,uid=0,loop=/dev/loop0  
-o loop,loop=/dev/loop0
```

Q Recently, due to a sudden power failure, I lost the entire data on my hard disk drive. Is there any software to recover lost data?

—R. Jagaraman, by e-mail

I would suggest that you take the help of some professionals to recover data, especially if it is very critical. Still, you can try a few of the open source applications mentioned below, as they do not have any restrictions like trial software.

- PhotoRec: cgsecurity.org/wiki/PhotoRec

- TestDisk: www.cgsecurity.org
- ddrescue: ddrescue.sourceforge.net

These tools can also be found on some Live Linux-based CDs (Knoppix, Ubuntu Rescue Remix, SystemRescueCD, etc). You can Google for information on how to download and use these tools.

Q I'm using CentOS on my desktop and want to know if there are any ways by which I can check my swap usage. I have 512 MB of RAM and have defined 1 GB as swap partition.


—Ashish Mitra, by e-mail

You can check swap usage by using any of the following commands:

```
cat /proc/meminfo
free -m
free -g
cat /etc/passwd
```

Q I have been a Windows user for a long time and used FileZilla as my FTP client. Currently I am working on a remote Linux computer with SSH access. Please guide me on how to download a file from another remote FTP server to this server.

—Prakash M., Madurai

You can use the command line *ftp* tool available. To download the file from another remote FTP server, you need to first log in to the FTP server by typing the *ftp <host name>*. This will prompt you for your username and password, if required. After logging in, you can use the *get* or *mget* command to download the file to your local server. To exit the *ftp* prompt, just type *bye*. Check out the man pages about this command for more details. 

Which one is your career **goal?**



medical



engineering



computer software



computer hardware



hotel management



clinical research

We help you **decide**

Hundreds of Educational Institutions Offering
Professional Courses & Training Under A Single Roof

EduTech's Highlights :

- * An event totally focused on professional education
- * A one-stop facility for students seeking admission into professional courses
- * Counselling by industry top-shots (HR and recruitments)
- * Major plan for advertising and promotion of EduTech Expo 2009
- * Powered by EFY Group's vision to boost professional education
- * Courses catering school pass-outs & fresh graduates

**EDUTECH
EXPO
2009**

13th - 14th JUNE 2009
PRAGATI MAIDAN, NEW DELHI



EFY Enterprises Pvt Ltd,
D-87/1, Disha Ind. Area, Phase-I, New Delhi-110020;
Ph: 011-26810601-03, Fax: 011-26817563



Dell Has Everything to Beat Apple



Does the PC manufacturer have the will to take the controls back from Apple, or will it continue losing market share to the Mac?

*W*indows is passé! Nobody seems to talk about it any more, except for, well, Microsoft! And even the folks from Redmond may shy away if you say something that sounds like 'Vista'. Now, the rage seems to be the Mac, and GNU/Linux, of course.

Apple's Mac has increased its share in the market post the Vista release. The market was ready for a new OS, but Vista disappointed users. GNU/Linux is putting up a good fight, but it has some distance to cover with respect to a few niche desktop use-cases.

Apple's heat

The heat created by Apple is burning the fingers of players like Dell, HP and many others. Owing to the soggy operating system that runs on their machines, they are losing market share to Apple. They have no other option but Vista, which simply turns users off and thus forces them to switch to the alternative called Mac. Dell did try to fight back by ceding to customer demands to continue shipping with XP, but you can't win a race in a vintage car when Apple is rolling out Ferraris, right?

The end user is in an even sorer situation: the lack of an alternative option by PC vendors forces them towards Macs, and the moment they buy one, they submit themselves to be locked in Apple's airtight compartment, where every bit of freedom is denied and every breath you take is controlled by the company. On top of it, Macs are expensive too. Well, so is Vista. The reason people go for Macs is the 'cool' factor. The graphics, the style, the looks, and the interface—everything is up to the mark, almost perfect.

So, those who want an awesome experience, choose Macs and those who can't afford it, or are simply willing to compromise, go for Windows. Of course, the smarter lot goes for GNU/Linux.

No doubt, GNU/Linux is a killer in the server space and in many other domains like embedded systems. An

ordinary home user can do a whole lot on GNU/Linux, but not everything. Professionals can't use high-end software for tasks like film, audio and image editing. Yes, there is the GIMP, but no Photoshop.

This is less of a problem for Microsoft, but a pain for vendors who have a very strong focus on desktop users—think of Dell! The company recently launched a much-hyped series of computers, branded Adamo. It looks good... it is slim. Well OK... but what does it run on? Umm... Windows Vista. "No, I don't want to spend that much money to buy another Vista. I will go for a Mac this time!"

Well, you can see Dell sweating! But, Dell is one company that can rewrite history and become the company with the highest selling machines in the world. The company now has the muscle and the curves—look at Adamo and you will know what I mean. It just misses one core thing, and that could turn people off.

Dell can do an Apple

Dell can have its own OS running on its machines. An OS that will be much more secure, and that will beat Mac in design and usage. What's more, Dell doesn't have to do much—all it takes is the desire to rule the market.

The OS is already there, you know what I am talking about—GNU/Linux. Dell can use Ubuntu, one of the most popular GNU/Linux distros, or even Debian... or anything else, for that matter. All it needs to do is customise the flavour of GNU/Linux to a level where it gets the glossy looks and glitz of the Mac—heck, it is possible to outshine a Mac.

Top priority

Dell can work with teams of developers to give a cool look and feel to the GNU/Linux operating systems. Check out the ease of GNOME and the look and feel of KDE, and you've got everything you'd want as a home user. Macs could be looking for a place to hide if Dell

could combine the best of GNOME and KDE.

From the usability point of view, GNU/Linux is not hard to use. Those who have used it recently would agree. Why, now you even have pre-packaged software just like .exe files, which you can click and install. Go to GetDeb and you will see what I am talking about.

Next in line

The second thing that Dell should do is tie up with players like Adobe and the rest to provide binaries for high-end software like CS4, et al, to fill in the remaining gaps.

The parallel movement to open source everything can continue without having GNU/Linux users whine or compromise when it is about something they have to complete within this lifetime. For example, if I have to edit movies, I can't wait till I'm 40 for some efficient free software that can do the job.

If there are binaries of the desired programs, I can buy the licence and use it on a GNU/Linux machine. In which case, even hard-core media professionals will be rushing to buy Dell machines running GNU/Linux instead of Macs. Or, at least, there will be some healthy competition.

Support

There is a strong GNU/Linux community out there, so Dell won't have to worry much about support. You will probably get better community support for GNU/Linux, compared to Windows or Mac support—google for it and you will know. Home users rarely buy support, and then with a strong community, much of Dell's problems are already solved. As for some minor issues on GNU/Linux machines relating to device driver support, since they're Dell's machines, all hardware can be supported. In fact, Dell can optimise the operating system so as to get the best out of the hardware, whether it be longer battery life for laptops or a much better overall user experience.

For corporate users, Dell can tie up with open source vendors to offer its customers various levels of support.

This will also create a healthy market for support which hitherto is being controlled by a few.

Invest money in the right place

The money Dell pays to Microsoft for the Windows licence can be used in customising Ubuntu. And since Ubuntu has a new development cycle of six months, Ubuntu will sport new features and looks more frequently than the Mac—do I need to even mention Windows here?

Don't imitate a bad Apple

There's one caveat for Dell. Don't make anything proprietary. Except for the binaries for non-free stuff like CS4, Ubuntu must release the code for everything they improve upon so that the community keeps on building on top of it, and Ubuntu keeps on improving, in the process.

This will enable Dell to create a unique proposition by getting extremely popular among GNU/Linux users, since sometimes, despite all the philosophy about freedom, Apple's

Mac does tempt them. But with Ubuntu, they will be able to showcase their shiny machines too!

Going forward

Gradually, with Dell becoming the most secure and popular GLPC (GNU/Linux Personal Computer), all the independent hardware manufacturers will also start rolling out open source software and drivers for their tiny whiny devices, so that they can run on the world's most popular operating system.

Those players will soon learn the beauty of open source—even if you have some bugs and less technical capabilities, the huge community is more than willing to take care of it, sooner than later.

This would lead to a phenomenon of releasing source code for hitherto proprietary technologies. Except for a few cases, it can also drive even the hard-core proprietary companies to start releasing specifications and code for their products.

The fear factor

Releasing code doesn't mean competitors will steal the show. Look at KDE and GNOME, Joomla and Drupal, or OpenOffice.org and the rest, as well as MySQL and PostgreSQL—despite the code being available, it's not that one is copying the other. Each has its own merits and they operate in a healthy competitive environment.


Then, since RMS is around with GPL, no one need worry about code being stolen and used, except for BSD-like stuff where Apple steals a lot from and returns almost zilch. Ironically, it charges heavily for what was initially developed by the community, and then puts such heavy locks on it that no one can peep inside. Don't forget, for every little problem, you have to visit an Apple Store and pay heavily.

It's the right time, Dell!

This seems to be the right time for Dell. Vista is taking its last breath. PC market share is falling. Should Dell wait and put its fate in the hands of another company and have no say in the software being used on its machines, or should it take control of everything that runs on its machines and become the No. 1 player in the world?

Dell can create history. All the other players, like HP, Lenovo, Asus and Acer, lack the extra muscle Dell has—customisation of hardware, the look and feel, and the desire to address the needs of end users. I don't think they have platforms like Dell Ideastorm, where they listen to what customers want.

Dell can do a lot. It can be the next 'sweet' Apple. But, the question is: can Dell hear the door bell?

The clock is ticking and Apple is closing in. 

By: Swapnil Bhartiya

A Free Software fund-a-mental-ist and Charles Bukowski fan. Swapnil also writes fiction and tries to find cracks in a proprietary company's 'paper armours'. He is a big movie buff, and prefers listening to music at such loud volumes that he's gone partially deaf when it comes to identifying anything positive about proprietary companies. Oh, and he is also the assistant editor of EFYTimes.com.

**Life as an administrator
is complicated enough...**

**But Backup and Recovery for your
Linux Servers does not need to be.**

NetVault: Backup simplifies backup and recovery
without compromising **Functionality** and **Scalability**

NetVault: Backup provides unmatched Data Protection for all major variants of Linux. We are now offering you the chance to see just how good NetVault is at no cost.

We provide continuous data protection (CDP) for your Linux servers and advanced application protection and recovery for MySQL, PostgreSQL, Oracle, Exchange, DB2 to name just a few.

Permanent FREE Use Edition for Linux is available for download at <http://www.bakbone.com/nvbu/redhat/>



NetVault is a true Linux data protection solution featuring:

- ☐ Online backup
- ☐ Point and click recovery
- ☐ Fully automated protection and recovery support for Linux based applications
- ☐ Virtual Tape Library (VTL) with disk staging
- ☐ SAN support with LAN free backups.
- ☐ Backup duplication for off-site storage of backups
- ☐ Support for MySQL native replication, restore to table level and to alternative databases or instances.

BakBone
The Power of Simplicity™

For more information, please contact:
✉ : IndiaSales@bakbone.com
☎ : +91-11-42235156



Debian Lenny

Brewed to Perfection

After some delay, Debian 5.0—dubbed Lenny—was finally released this Valentine's Day. How appropriate! But does it really lead to true love? There is only one way to find out...

I'm a big fan of Debian and the dedicated community involved in developing it. I make no secret of that. When I reviewed Etch (4.0) last year, I declared that if I were to finally grow up and settle down with just one distro, *this* would be the one. I like the fact that it's not backed by any commercial entity and sticks closely to its Free Software principles.

Vital stats

- Distro base: Itself; Debian is the mother of a lot of modern distros
- Packaging: *.deb* (managed by the mighty APT)
- Linux kernel: 2.6.26
- Default desktop: GNOME 2.22.1

Installation

I downloaded the Lenny network install CD image, my preferred installation method with Debian. I find it a bit wasteful to burn five full DVDs and then only use one of them. There's some merit in having everything downloaded before you start, but if you've got a reasonably fast and stable Internet connection, I recommend the network install personally. I was fortunate enough to have a chat with Steve McIntyre (the current Debian Project Leader) on a recent episode [linuxoutlaws.com/podcast/77] of Linux Outlaws, during which he told us that there are even Blu-ray images available now with everything on one disc. That sounds amazing, but not having the small fortune it requires to purchase a Blu-ray burner at the moment, it wasn't an option for me, sadly.

I booted up from the CD to get started, as always, testing it with my trusty Dell XPS m1330n laptop. I chose the graphical install and proceeded through the usual prompts for language, location and so on. On the surface, the installer looked the same to me as the Etch installer—I'd heard people raving about 'the new Debian installer', but I don't see any evidence of it. It's not necessarily a bad thing, it works well and as the old adage goes, "If it ain't broke, don't fix it."

I was informed pretty quickly that there were no drivers available for my Intel wireless card, but was instead offered the option to load them from a disc; I didn't have one. It wasn't a major problem since I could use the Ethernet port for the network install anyway. The base system was then installed from the CD really quickly... I mean, *lightning* quick.

All the additional packages need to be downloaded from a network mirror and you're offered the chance to pick one near you. Picking the one nearest to your location will usually result in much faster download times, so choose carefully. I picked the default UK server and was prompted to choose between installing the following package sets: Web server, mail server, DNS server, desktop, laptop and many more. I chose the desktop, laptop and standard groups, leaving all the server options alone.

The installer then informed me that I was in for a wait while it downloaded 817 packages and installed them. After this, GRUB was configured and I was ready to reboot. The whole process from putting the disc in, to booting up my new desktop, took approximately 25 minutes. That's

a respectable period of time when you consider a large chunk of that was consumed by actually downloading the software. If you use a full install disc, it would probably be a lot quicker.

It was a pretty quick and simple install, so I think most users could handle it without breaking into a sweat.

Configuring the system

Now, this being Debian, it's a little different from the likes of an Ubuntu or openSUSE, which try to configure every little thing for you automatically. Yet I still had a lot of configuration to do. The first order of business was to get my wireless card working, which actually turned out to be a lot easier than I'd feared. I opened up Synaptic (the graphical package management tool) and then looked at the preferences to see what software sources were enabled.

I could see a couple of groups weren't ticked, such as software that doesn't comply with the Debian Free Software Guidelines—it seems the drivers for my wireless card fall into this group. I enabled non-DFSG software, refreshed the package list and searched for 'Intel', which yielded a few packages, but scrolling through the list I quickly found one relating to my Pro Wireless 3945 card. After a minute, the wireless light started flashing and I was able to log into my home network easily. Not too difficult, but most distributions have this driver working out of the box. It's only omitted here because it doesn't satisfy the DFSG that I understand.

I have to say at this point I was really impressed by the speed and responsiveness of the system—it seemed very snappy and a quick look at the resources monitor showed I was only using 145 MB of RAM, even with quite a few large programs open on my GNOME desktop.

Most of the basic software you would need is included by default—like OpenOffice.org, The GIMP, Rhythmbox, Pidgin and Iceweasel, which is great. Iceweasel is essentially a clone of Mozilla's Firefox, created in protest against the actions of Mozilla Corporation and its licensing process. It's actually repackaged from Firefox source code with minimal modifications, so all your extensions will work just fine. I've never had any problems with Iceweasel or noticed any real difference from Firefox to speak of.

I installed most of the additional software I needed via Synaptic and could get packages for other things like Gpodder [gpodder.org] from the developer's own website. Installing the .deb files wasn't quite as simple as double-clicking them, though. I had to open a terminal and use the 'dpkg -f' command, but I got there. I also noticed that the *Add/Remove Software* tool from Ubuntu is now included.

There were a couple of programs I couldn't find in any of the repositories, even the experimental ones, Gwibber [launchpad.net/gwibber] being a prime example. I had to download all the dependencies for it manually and then install from the source code I'd checked out with Bazaar [bazaar-vcs.org]. It was a bit awkward and at times I felt I wasn't doing things 'the Debian way' by trying to install cutting-edge software. The distribution is renowned for its stability and some critics would say, its slow adoption of



Figure 1: Initial desktop



Figure 2: Iceweasel with Swifdec

new software. That's a trade-off for the stability, I suppose.

The freedom-hating drivers for my nVidia graphics card were installed easily with Synaptic. I then had to run the *nvidia-xconfig* command manually in a terminal to update the X server settings. I installed the *nvidia-settings* package because I use it for switching between internal and external displays.

I also like to use the Compiz Fusion 3D desktop. Some people dismiss it as eye candy or frivolous decoration, but it's far more than that. The easiest way to install Compiz on Debian is through the fantastically named Repository Of Shame [shame.tuxfamily.org/repo]. It contains everything you could need and provides handy meta packages for easy installation.

I should take a moment to point out the utterly tragic news that Shane Lee (a.k.a. Shame), maintainer of this repository and well-respected Debian developer, sadly died last November. It's hard to know what to say. Words fail me but he was only 35 years old and he will be sorely missed. My deepest condolences to his family. The repo still seems to be working and I was able to install everything by following the simple instructions on the site.

It's hot, damn hot!

This being the stable branch of Debian, a lot of the software is quite old. Some prefer the term mature, but I missed things like GNOME 2.24 with its tabbed Nautilus file browser and other improvements. I looked at adding the latest GNOME, but it



Figure 3: My finished desktop

didn't seem like an easy process. There are some guides on how to compile and integrate it, but I decided to pass that. In many ways it would also have cancelled the main benefit of running Debian anyway—tried and trusted software that you know will work. If the Debian guys say something is stable, you can believe it's stable—a conservative approach, but I admire their commitment to quality.

Lenny also ships with a much older kernel than I would normally use. My machine kept overheating quite badly, and the fan sounded like a small aircraft attempting to take off or perhaps like trying to watch a film on an Xbox :-). There were some issues with my model of nVidia card soon after its release, till modifications were made to heat management in the Linux kernel to compensate. It now works great and the fan does not normally need to work overtime with the newer kernel. Consistently though, with Lenny I found I needed to leave the machine to cool down. This may be my fault for having the evil nVidia card, but it definitely didn't like the older kernel.

Conclusions

- Ease of installation and use: 2/5
- Stability: 5/5
- Community and documentation: 5/5
- Features: 3/5
- Overall: 4/5


I was able to get the system working as I wanted, under Lenny, but it took some time and expertise. I compiled a lot of the newer software from source when I couldn't find packages. I did try adding the Squeeze repos [www.debian.org/volatile] and even some experimental ones, but still found a lot of the packages were significantly behind. That's what you expect with Debian though—you don't use it because you want the latest cutting edge software; you use it because you want stability, speed and the support of some of the best developers in the business.

There's a very useful guide [www.howtoforge.com/the-perfect-desktop-debian-lenny] to setting up the perfect desktop with Lenny and it introduced me to the *Debian-multimedia.org* repository. I highly recommend it to everyone to easily add extra multimedia programs.

The overheating problems due to the older kernel were a shame. I could of course install another kernel or even compile one from scratch, but for me, that seems to defeat the purpose of using Debian. It would be easier to use Arch or another distro if that's the aim. Back when I tried Etch, it was on a desktop machine and I think the main difference this time around is I'm on a notebook. I said while reviewing Etch that I was amazed at how much of Ubuntu is Debian painted brown, and that's still true. But I also gained a new appreciation for the work the Ubuntu guys do in making their software run effortlessly on laptops. I like that Debian is very true to its Free Software ethos, including things like Swfdec instead of the proprietary Adobe Flash player. However, in practice I found it didn't work very well, and I ended up ripping it out anyway and installing the evil Adobe Flash.

Do I still love Debian? Yes, of course, I do. And there's no doubt about that. However, I have to confess that maybe I've changed. On my laptop I want the latest and greatest software to play around with. I like to get the newest GNOME, for example, and Debian just isn't the right distro for that. It does have many strengths—Apt is still Debian's killer app for me and I've yet to find another package manager that really compares.

The bottom line though, is that on the desktop I find Debian too conservative for my tastes these days—it's a personal thing. On a server or even a desktop with a specific purpose, it excels and I wouldn't replace it with anything else. If I need a Web server, database server or... well... any other kind of server, I'll be reaching for my Debian CD, because I know it will be rock solid. But on my laptop, I guess I'll keep hopping around in search of new software.

If you've used Linux before and you're not afraid to learn a little bit about its finer workings, Debian has a massive amount to offer and I'd encourage everyone to try it at least once in their life. It's not for everyone, but it's still a great testament to what a group of dedicated volunteer developers can achieve on our behalves, when they come together. Try it out for yourself and let me know what you think. 

Resources

- Debian's home: www.debian.org
- Full installation slide show: www.flickr.com/photos/mvmetri/sets/72157615895122780/show
- Debian Free Software Guidelines: www.debian.org/social_contract#guidelines
- Useful guide: www.howtoforge.com/the-perfect-desktop-debian-lenny

By: Dan Lynch

Dan is a 28 year old writer, musician (shedmusic.net), broadcaster, producer and convicted geek from Merseyside in the UK. He is mad about all things open source, particularly GNU/Linux, but also BSD, OpenSolaris and various Web/mobile technologies. He does a bit of programming and Web development, but hasn't done any really meaty projects in a while. You can find out more about him at danlynch.org.

This work is released under the Creative Commons Attribution-Share Alike 2.0 UK/England & Wales license (creativecommons.org/licenses/by-sa/2.0/uk/), and was first published at danlynch.org/blog/2009/03/lenny.

Microsoft

EQUIP YOURSELF

Microsoft invites all developers to join the open source community and online resource center. Interact with the likes of Sam Rami, Director of Open Source Software Ltd, Jonke Ewing, Lead of OSS Community & Platforms and many others to find solutions to all interoperability challenges.

JOIN WWW.PORT25.IN, A MICROSOFT PLATFORM FOR OPEN SOURCE DEVELOPERS.

PORT25



McGraw-Hill



Solaris

for your Desktop

OpenSolaris 2008.11 is the latest release of the OpenSolaris operating system. In this article we take a look at what it has to offer, besides all that's new.

The OpenSolaris 2008.11 live CD boots to a minimal OS. You can explore the OpenSolaris OS on the CD without affecting the existing operating system on your laptop.



Note: If you already have OpenSolaris 2008.05 installed on your system, you can upgrade to 2008.11 without using the DVD by following the instructions here: tinyurl.com/opensolaris-upgrade-200811

Get it installed

OpenSolaris 2008.11 requires a minimum of 512 MB of RAM and a primary partition with a minimum of 3.1 GB of disk space (though at least 9 GB is recommended). If you do not already have an unused primary partition, use a tool like GParted to create one for OpenSolaris. You can set the partition type to 'Linux SWAP' for the time being. Make sure you back up any important data on the machine before resizing/creating partitions.

IMPORTANT: If you are installing OpenSolaris on a system that already has Linux installed, remember to boot into Linux and back up the `/boot/grub/menu.lst` file—copy it to a thumb drive, mail it to your GMail account

or even just write it down on a piece of paper. OpenSolaris overwrites the installed GRUB with its own version during the installation. You will need to add the entries you had in your Linux `menu.lst` file after installation.

Insert the OpenSolaris live CD into the computer and set up the BIOS to boot from the CD drive, then boot into the live CD. Choose the first boot entry in the boot menu. After a while, it will go to the console mode and ask you to choose your keyboard layout. Defaults will suffice in most cases. In a few minutes, you should be looking at the blue GNOME desktop (Figure 1).

This release includes much better hardware support compared to previous releases. You can use the Device Driver Utility (DDU) on the live CD to check if any of the hardware on your system is not supported (Figure 2). [After installation, you can find the DDU under *Applications*→*System Tools*→*Device Driver Utility*.] If everything is fine, you're ready to take it for a ride... err, I mean, install it on your hard disk.

Double click the 'Install OpenSolaris' icon from the live desktop. (If you are prompted at any point to log into the live CD, both the user name and password are 'jack'. The root



Figure 1: The default OpenSolaris 2008.11 GNOME desktop

password of the live CD is 'opensolaris'.)

The first step in the installation is choosing the partition to install OpenSolaris into (see Figure 3). Unless you want to blow away the data on the disk and use it all for OpenSolaris, make sure the 'Partition the disk' option is selected. Set the partition type as 'Solaris' for the partition you want to use. There must be exactly one partition of type 'Solaris' for the installation to proceed. (If you already had a Linux swap partition configured, it might show up as another 'Solaris' partition. Change this to 'Unused' to make sure the installation does not touch this partition.) Reconfirm that the partition chosen is the right one and click on *Next* to continue.

In subsequent screens, you can configure the time zone, date, time and locale. The final step is to set up the user accounts. A root login is not enabled on the installed system and it is not mandatory (though recommended) to enter a root password. Make sure you create a user account for yourself.

Click *Next* and review the installation settings (Figure 4). Use the *Back* button to go back to previous screens if you need to change any of the settings, else click *Install* to start the installation process. Figure 5 shows the installation in progress.

In your hard drive

Once the installation is complete, the *Reboot* button will boot into OpenSolaris. Remember to remove the live CD from the drive and preferably also reset the boot order in the BIOS to pick the hard drive first. Once the OpenSolaris login screen comes up, log in with the user name and password you have set up.

If you have trouble logging in after a reboot, look at the 'Troubleshooting OpenSolaris logins' link in the *References* section at the end of the article.

Restoring Linux GRUB entries: This step is not necessary unless you had backed up the *menu.lst* file from your existing Linux installation. The 'bootadm list-menu' command will display the location of the active GRUB



Figure 2: The Device Driver Utility program



Figure 3: Where do you want to install the OS?



Figure 4: Installation summary screen



Figure 5: Installation in progress




Figure 6: OpenSolaris Package Manager

menu file. Edit this file and add the entries from your Linux *menu.lst*.

Root shell: The first user to be created has super user privileges. That user can prefix any command with *pfexec* to run it with super user privileges (similar to *sudo* the command Linux users are familiar with). For example, to get a root shell, run *pfexec bash*.

You are now all set to start exploring your new OpenSolaris installation! The first thing is the customised default artwork—I don't know what others think, but I find it quite attractive.

OpenSolaris 2008.11 uses GNOME 2.24 as the desktop and introduces a whole bunch of improvements, including an improved package manager, automatic backup management, and a desktop search tool, among others.

More software

The package manager can be used to install software from a collection of network repositories (Figure 6). Software available on the official OpenSolaris repositories

include OpenOffice.org, Netbeans, Eclipse, Songbird, Gobby and the optimised OpenSolaris Web Stack including frameworks like Drupal and Django, several of them instrumented with DTrace probes allowing you to dynamically trace your Web application for improved performance. (You can get more details about the Web Stack at tinyurl.com/sun-web-stack.)

You can launch the package manager from the toolbar or *System→Administration→Package Manager*. The interface should be familiar to anyone who has used Synaptic or one of the other GUI package managers in Linux.

Note: If you prefer the command line, run *pkg help* in your terminal window to get details of the command line interface.

Some encumbered software that cannot be freely redistributed because of licensing constraints are available in the OpenSolaris Extras repository. Access to this repository requires you to sign up for a Sun Online account. Point your browser to pkg.sun.com and follow the instructions to sign up and log in. You will be redirected to the 'Certificate Requests' page where you can request access to the OpenSolaris Extras repository. Follow the instructions on the page to set up access to the repository. (You will need to run a bunch of commands on the command line. The package manager GUI does not support this currently. However, once you have completed the one-time task of configuring access to the Extras repository, you can use the GUI to browse and install applications.)

Once you have the Extras repository configured, you can install packages from there—for instance, VirtualBox or the Flash plug-in for Firefox. Make sure you choose the right repository in the top panel of the package manager before you start looking for these packages.

The package manager now also includes an update manager that notifies the user when there are updates available.

What's in a filesystem?

OpenSolaris uses ZFS as the default filesystem. ZFS is a new kind of filesystem that provides simple administration, transactional semantics, end-to-end data integrity, and immense scalability. Unlike conventional filesystems, ZFS has integrated volume management features that make it much easier to perform traditionally complicated tasks like resizing filesystems, adding new disks, etc. For example, adding a new disk is as simple as adding the device to the ZFS storage pool (*zpool add*). The extra space is instantly available to all filesystems in the pool.

ZFS also has the ability to create instantaneous snapshots and clones of your filesystems. Snapshots and clones occupy no additional disk space in the pool when initially created, and any unchanged data is shared among the filesystem and all its snapshots and clones.



Figure 7: The Time Slider icon in the Nautilus toolbar



Figure 8: Tracker Search Tool

One of the places that the snapshot feature of ZFS is on display is the new Time Slider.

What in the world is a Time Slider?

The Time Slider provides a way to set up automatic backups of your data on the same disk, enabling you to browse and restore older versions of files using GNOME's Nautilus file browser. Anyone who has ever accidentally deleted or overwritten an important file will appreciate the time and effort saved by being able to walk back in time and pick up an older version of the file without even leaving the file browser!

Time Slider is disabled by default. You can enable it via *System→Administration→Time Slider Setup*. To browse your backups, open the file manager, navigate to the directory you are interested in and click the icon with a clock (Figure 7) in the Nautilus toolbar ("Browse the current location snapshot history"), and use the slider that is displayed to go back in time to a view of the directory from the past. To restore an item, right-click on it and choose 'Restore to Desktop'.

Can't remember where a file is... now what?

Tracker is a desktop search tool used to locate data from the system, including text documents, music, video/image files and e-mails. It indexes the filesystem in real time and automatically keeps its metadata up to date.

Tracker is not enabled by default. To enable it, go to *System→Preferences→Sessions* and enable the "Tracker" and "Tracker Applet" start-up programs. You might also want the tracker to index your home directory. To enable this and other configuration tweaks, start up a terminal window and run the 'tracker-preferences' command.

You can launch the tracker search window by clicking the magnifying glass icon in the toolbar.

Network auto-magic

Several enhancements have been made to the simplified network management infrastructure (NWAM) in OpenSolaris 2008.11. This, coupled with a vastly improved set of wireless drivers supporting WPA, ensures that OpenSolaris 2008.11 gets closer to, "it just works!"


Right-click on the network manager icon in the system tray to switch between wired and wireless networks, and to choose from among the list of available wireless options.

What else?

A new Print Manager is included in the release, which provides support for the management of printers, auto-detection of USB printers, etc. CUPS support is planned for a future release.

OpenSolaris 2008.11 includes suspend to RAM support for a limited set of laptops from Toshiba (Tecra M8, Tecra M9, Tecra M10, Tecra A9, Satellite A205, Portege R500), Lenovo (6460SVU, 889703U) and Dell (Latitude D630, Precision M4300), apart from all Sun Ultra 20, Ultra 24 and Ultra 40 workstations.

Finally, a lot of utilities that people moving from Linux expect to find are available by default, though there are OpenSolaris alternatives available, which are preferable in most cases. For example, *sudo*, *top* and *slocate* are installed by default—the OpenSolaris alternatives being *pfexec*, *prstat* and *tracker-search*, respectively.

If you have any issues installing or using the live CD, head over to the release notes wiki at tinyurl.com/lfy-opensolaris-release-notes and check if the problem has already been reported. If not, please add the details in the 'Discussion' tab on the wiki page. Also, you are welcome to join the Bangalore OpenSolaris User Group at opensolaris.org/os/project/bosug where you can discuss these and related issues. 

References

- OpenSolaris 2008.11 Release Notes Wiki: tinyurl.com/lfy-opensolaris-release-notes
- What's New in OpenSolaris 2008.11: www.opensolaris.com/learn/features/whats-new/200811
- Upgrading from OpenSolaris 2008.06 to 2008.11: tinyurl.com/opensolaris-upgrade-200811
- Troubleshooting OpenSolaris logins: tinyurl.com/opensolaris-troubleshoot-login
- Sun Web Stack: tinyurl.com/sun-web-stack
- Bangalore OpenSolaris User Group: opensolaris.org/os/project/bosug

By: Venky TV

Venky has been working at Sun for the last 9 years, is one of the founders of the Bangalore OpenSolaris User Group and believes very strongly in the virtues of laziness.

Meet the Messiah

March 12-14, 2009. That was Open Source India 2009, Tech Days. Here's a report for those who couldn't be with us at hot, yet happening Chennai.

"It was the best of times, it was the worst of times..." these words borrowed from Charles Dickens' novel *A Tale of Two Cities* seem apt today. The world is facing an economic slowdown. Businesses are shutting

down, or at best downsizing. At a time like this, we certainly need a messiah to save us, or at least offer us some relief so that we can continue to survive (and, more importantly, grow). And we found such a saviour in open source.

Today, no company can execute its operations without a significant amount of IT—it's pretty much the backbone for any business. And in hard times like these, maintaining and scaling the IT infrastructure requires investments. However, not many would like to touch the reserves they have. So, what should they do? Open source seems to have the answers.

We did not have this in mind during the initial phases of planning our annual open source business conference. However, once it was clear that we had to address the need of the hour and try to help the industry find some answers, we started looking for experts in the field to share their experiences and *that* shaped the Chennai Chapter of Open Source India 2009—Tech Days.

We realised there were two areas we needed to touch upon. One was the 'soul' of the open source industry—the community and the developers. The second was the enterprise segment, made up of the corporate users. The

Panels, panels and more panels!

Before this edition of Open Source India, we thought handling multiple speakers and tracks was a challenge. But, right on the first day of OSI 2009, we knew that we had taken up an even bigger challenge by planning most of our sessions as panel discussions! We always knew that shifting to a panel discussion format would mean working with many, many more speakers, but what we had not anticipated was the complications that needed to be handled back stage.

Have all speakers reached the venue? Have all of them been introduced to each other? Are they all in sync with the goal of the panel? Phew! Every panel seemed like handling a whole day's content. But, going by the feedback we have received online from many of you—it seems the effort was worth it. Through the panels, we had hoped that the sessions would be more interactive and the interest kept alive amongst the audience—even in the post-lunch sessions! Plus, we hoped that the audience would be able to cross-question the speakers and get a wider and more unbiased perspective. Thankfully, it seems like they did. So, will we have panels in the next edition, too? Most certainly.

There was a Miss South India contest scheduled in CTC the day we landed in Chennai. While our editor, IT admin and logistics guy blissfully bought off the peesses and indulged in sin, our correspondents worked frantically to sort out last minute issues back at the hotel!

content of the event was developed with these two focus areas in mind. To address the needs of the industry, we created an IT Managers' Track. And to connect the new generation of developers (or those unaware of the advantages of the open source development stack) to the community, we created the Software Developers' Track.

First day, first show!

Steve Lau, senior consultant of HK OSS Centre, did the honours at the lamp lighting ceremony, along with Dr P. Balasubramanian, head of open technology centre, National Informatics Centre (NIC), Dr S. Ramakrishnan, director general, C-DAC, Dr M.R. Rajagopalan, executive director, C-DAC, and Poonam Kapoor, director, EFY Group.

Following this, Dr Ramakrishnan delivered the inaugural address. He focused on the impact of free/libre and open source software all over the world and India's role in spreading it. He said that India had immense manpower in open source, and if we used it fruitfully it could be a tremendous power, worldwide.

"The key thing in spreading open source is to adopt it. Issues like interoperability and LSI compliance in open source are being sorted out. Here the government is very supportive to the idea of FOSS and many governmental projects are being done using it," he conveyed. He also added that there is tremendous opportunity in the field; however, it still requires a lot more community-based initiatives and forums to groom it.

Steve Lau delivered the keynote address. He pointed out that the Asian OSS Centre, which covers many countries—including India, Hong Kong, Pakistan, Japan and Indonesia—was set up on March 27, 2008 to promote FOSS in Asia. Its purpose is to facilitate the information and knowledge exchange

among the participating countries. "Asia has huge opportunities for open source and there is wide scope for learning in the field. Open source will give immense job opportunities to entrepreneurs, and that's why Asian countries should introduce and promote open source in various projects," he elaborated. Dr Balasubramanian also spoke during the inaugural session.

Post the inaugural addresses, the event was split between the two parallel tracks, one for IT managers and decision-makers, and the other one targeting developers.

Software Development Track

This track was aptly titled FOSS-powered Software Development. It witnessed discussions on various topics ranging from software development technologies, FOSS databases, tools for creating top quality software, etc.

The session was kicked off by Dhiraaj Sinha, head of technology at Perot Systems' Applications Solutions Group. He stressed on where and why Perot Systems is using open source technologies for its software development requirements. This was followed by a panel discussion that guided the audience on how to migrate from the proprietary to the FOSS stack for development.

During the course of the day, P.K. Mishra, principal consultant architect, Novell, said that users have a variety of choices available in FOSS compared to the small footprint of Microsoft. Pointing out Mono, he noted that the platform is a complete FOSS implementation of the .NET framework, which runs on both Linux and Windows, amongst other systems. He added that the .NET apps that are meant for compilation in Visual Studio can be compiled without any changes directly on Mono, which means you have your favourite .NET applications running on Linux.

Supreet Sethi, who likes to describe himself as a photographer by choice (although a freelance software developer by profession, just so he can afford his passion for photography), took up the topic of FOSS databases. He pointed out that many of the conventional

databases worked almost the same way. A vendor generally recommends a particular database for a specific use. For example, MySQL is recommended for content management, yet if you are in a transaction process company, MySQL is the worst option.

John Smedley, senior architect, Ingres Corporation, elaborated, "If we are looking for a specific product to solve a specific problem, we need to know the problem thoroughly. We need to know the needs to build a prototype according to it."

Ashish K. Mathur, software architect, IBM, stressed on the importance of 'testing' in open source. He also turned the spotlight on the unique spirit of volunteerism that drives innovation in the open source world. "Nobody is doing it [testing software] because they were asked to do so," he said.

The importance of a version control system (VCS) to create top quality software was covered by Bhuvaneshwar A. of CollabNet. Technically, a VCS is a piece of software that tracks differences between every "Save" you do in your code (called 'commits' in VCS nomenclature) and stores it, so that it can recall the state of the file before any of the "Save" operations that you performed. This can be a great help in troubleshooting. If a bug was introduced into the code unknowingly and the issue was noticed, say, two years after the viral commit was made, with a VCS you could know when it got into the cycle and track it easily.

IT Managers' Track

Running parallel to the software development track were programmes

While it made our chests heave with pride witnessing over 50 speakers being eager to take up the podium and speak before the well packed audiences, it was also a pain to cut a few of our speakers short owing to shortage of time. During one such incidence, a speaker got royally pissed, dropped the microphone on the podium and walked off. Sir, we'd like to apologise for the inconvenience.

Linux Learning Centre

Pioneers in Training on Linux Technologies

Trained participants from over 36 Countries in 6 Continents

Linux OS Administration & Security Courses for Migration

LLC102: Essentials of Linux OS
LLC103: Linux System & Network Administration
LLC203: Linux Advanced Administration
LLC303: Linux System & Network Monitoring Tools
LLC403: Qmail Server Administration
LLC404: Postfix Server Administration
LLC405: Linux Firewall Solutions
LLC406: OpenLDAP Server Administration
LLC408: Samba Server Administration
LLC409: DNS Administration

Courses for Developers

LLC104: Linux Internals & Programming Essentials
LLC105: Programming with Qt
LLC106: Device Driver Programming on Linux
LLC107: Network Programming on Linux
LLC108: Bash Shell Scripting Essentials
LLC109: CVS on Linux
LLC204: MySQL on Linux
LLC205: Programming with PHP
LLC206: Programming with Perl
LLC207: Programming with Python
LLC208: PostgreSQL on Linux
LLC501: Programming with OpenGL

RHCE Certification Training

RH033: Red hat Linux Essentials
RH133: Red Hat Linux System Administration
RH253: Red Hat Linux Networking & Security Administration
RH300/301: Red Hat Rapid Track Certification Course

RHCSS Certification Training

RHS333: Red Hat Enterprise Security: Network Services
RH423: Red Hat Enterprise Directory Services & Authentication
RHS429: Red Hat Enterprise SELinux Policy Administration

**New Courses
RH142 & RH184**



**Fast Track Courses in April on
Linux Kernel Internals
Device Driver Development**

**RHCE Exams in Bangalore
on 9 & 24 April 2009**

**RH301 from 13 & 20 Apr '09
RHS333 from 4 Apr '09
RH423 from 11 Apr '09
RHS429 from 18 Apr '09**

**Linux Support & Solutions
Installation, Setup and Support Solutions
for RedHat, Ubuntu, SUSE, CentOS Servers**

For more info log on to:
www.linuxlearningcentre.com
Call: 9845057731 / 9343780054
Email: info@linuxlearningcentre.com

**RED HAT Training Partner
RHCE Examination Centre**



Linux Learning Centre Private Limited

Registered Office & Corporate Training Centre
635, 6th Main Road, (Adj. Bank of India) Hanumanthnagar, Bangalore 560019
Tel: +91.80.22428538 / 26600839 / 26610999 TelFax: +91.80.26600839
Cell: 9845057731 / 9343780054 Email: info@linuxlearningcentre.com
URL: www.linuxlearningcentre.com

LLC Satellite Centre - Bangalore
1291, 24th Cross, 30th Main,
BSK II Stage, Bangalore 560070
Tel: +91.80.26712928



▲ The panel shares its view on how to migrate to FOSS development stack



▲ Audience in the IT Managers' track



▲ Panel shares information on software development languages and platforms



▲ An attendee clarifies his doubt



▲ Boudhayan shares his reasons for fiddling with FOSS

addressed to the IT managers. The first session of this track was kick-started with the all-important question, "Why FOSS makes sense for your IT infrastructure?" The panellists for the session included Varad Gupta, head of Kean & Able, Dibya Prakash, technology manager at Azri Inc, Vivek Bhatnagar, country director and global alliances, Ingres Corporation, and Venkat Mangudi, CEO, Venkat Mangudi Consulting.

The session touched upon various aspects of IT infrastructure within an organisation. While Mangudi spoke about the typical IT infrastructure set-up in any firm, Gupta built a case of how, why and when to go for open source solutions. Dibya Prakash also covered the details on why any organisation needs open source, Bhatnagar shared with the audience the critical role database solutions play, and how open source solutions excel over their proprietary rivals. The session seemed to further bolster the audience's confidence in open source.

The session was followed by a panel discussion on enterprise applications. It was chaired by Mangudi, while the other panellists comprised Abhijit Das, director, Konsulte, D. Seetharam, government programs executive, IBM, Tim Cloonan, community manager, JasperSoft and Redhuan D. Oon, leader ADempiere (ERP) Project. The panellists shared with the audience the various solutions available and how they address different needs. Cloonan also pointed out the way people get attached to the community and thus help improve the product.

The third session of the day was another panel discussion on messaging and communication systems. This was moderated by Varad Gupta, and the panellists included Vikash Jha, project manager (Open Source) at HSBC Investsmart, and D. Seetharam.

Post lunch, 'network monitoring and management' was the topic for discussion. The session was moderated by B.C. Sekar, senior project manager, networking products division, HCL, and the panellists included Vikash Jha and Satyakaam Goswami, a freelance FOSS consultant. It was structured to be more like a Q&A session, where the audience threw questions to the panel. These doubts

We had a guest who wanted to showcase KDE 4.2 to the audience. While using one of our compere's laptops, he mistakenly deactivated 3 of his notes that he had on another Activity. Poor guy had to keep remembering his to-do list all night long.

ranged from the right kind of solution, to security in the network monitoring space. And the panellists gamely handled every question, putting to rest potential users' worries and misconceptions.

A double-edged discussion on security was the highlight of the evening's programme. PK. Mishra, principal consulting architect, Novell, represented the AppArmor side of the story, whereas Varad Gupta spoke up for SELinux. There was an interesting comparison between the two by supporters of the respective technologies.

The audience aired their concerns over certain open source solutions already deployed in their firms, as well as the problems they were facing with their current proprietary software-based infrastructure. While the panellists addressed these queries, they pretty much made it clear why open source is the way to go for an enterprise.

Day 2 Show 1: FOSS For Mobiles

Calling the response to these sessions a tsunami, would be an understatement! On the second day of the event, we were faced with a deluge of students, FOSS enthusiasts and IT managers. The hall hosting the IT Manager Track, dubbed as the *CXOSummit on Virtualisation*, was a full house, and we faced a situation where the rapt audience spilled on to the corridors in order to participate.

The situation at the Developers' Track was even worse, or should we say, better? The good news was that the hall with a capacity to seat 700 people was full. The bad news was that we had 500 more people to accommodate. The show was: *FOSS on Mobiles*. The actors were: Moblin, Android and OpenMoko. We were really elated at the success of the show, but we were completely at our wit's end on how to accommodate the 500 extra people. But did they mind? Nah, they were much

During one of our sessions in the developers' track, just when the last speaker stood up to present his ideas, we had an audio gone up and snippets of audio from the adjacent room streamed in with the announcement, "2 more minutes left". Dope!

The Tweetup taught us that it's all about 'free beer', and hardly about 'free speech' ;-)

A Tweetup is an event where people who use Twitter or Ident.ca come together to meet in person. Normally, we connect with our friends online after we have met them. At a Tweetup, you meet the people you might only otherwise know virtually.

LFY had very recently hopped on to the social media bandwagon and decided to make itself visible on the Twitterverse, as well. In the past two months, we managed to rake up more than a hundred followers. Given the kind of opportunity we got while organising an event as huge as OSGI '09, we couldn't resist meeting up, face-to-face, with a few of our friends and followers from Chennai. Of course, the fact we had beer on the house helped the cause.

We finally got face-to-face with those like @PlaneMad, @neranath, @indiblogger and @abattor, whom we'd only known online. While the agenda of the meet was to discuss the latest social media rage, businesses sprouting up around them and the marketing opportunities opening up, we wanted to figure out how to amalgamate

the open source ethos with micro-blogging services like Twitter and Ident.ca.

One of the most important, but probably not so startling, discoveries was that the social media world cares more about the work than the tool itself. As @PlaneMad had put it, "We do not really care whether a software is free or not, as long as it doesn't cost us a penny. Yes, we would always like to discourage piracy. But, that is the max we would care about."

This makes us wonder—if you are a marketer or a sales guy, would you really care whether a particular tool is free or not, as long as you stay within legal limits and shell out the minimum possible bucks?

Are you on Twitter too? Follow us @LinuxForYou.

After the tweetup, where we had free beer for one and for all, we still had a few beers left and had planned to collect the whole crate the next day. Out to next day, the whole crate was nowhere to be found.

too engrossed with the presentation to notice. And fuel was added to the fire by Narendra Bhandari, director of Asia Pacific, software and services group-developer relations division, Intel, who started off with a rocking video (*that* was super) on the screen, which got everyone hooked.

The hall managers were sweating, overwhelmed by the number of attendees that was beyond all expectations. Most were students, which indicates that the next generation is quite interested in the open source technologies. Bhandari discussed the evolution of smartphones. He drew an overall picture of the cell phone market that is constantly growing at a steady rate, stating that these devices are incorporated with operating systems that have custom interfaces to create interesting deviations of the user interface.

He said that most of these operating systems are Linux based. "We need cell phones or mobile data devices for entertainment, business critical uses, daily conveniences and a hybrid set of other uses. Moblin is a true open source platform for hybrid purpose devices. It is a Linux-based solution, which provides a rich media interface. You would want data on such devices to be synchronised with your home PC or to your corporate mail. Linux will be the viable choice in such cases," he stressed. For most of us who can't afford a Blackberry, *the times they are a-changin'* thanks to Linux-

powered mobile devices.

Rakshit Hooja, director, IDA Systems, while discussing "The OpenMoko Experience", said that smartphones have growing hardware capabilities that can be leveraged, and more and more simple development platforms will emerge in the coming days. "The OpenMoko platform consists of truly open hardware as well as software. The OM hardware can run multiple software stacks, such as the OpenMoko stack, Nokia QT Extended, Debian and so on," he confirmed. Hooja suggested that by using OM, your phone could 'grow'. Rather than buying a phone every two years to get the latest set of features, all you need to do is update the phone's software like you update the OS in your PC.

Next up was Thilak Kumar, technical account manager, engineering division, Wind River. His talk included an enlightening presentation on the architecture of the Android stack, and its various licensing divisions. We got to know that the kernel has the GPL, but the rest of the stack is supported by APL (Apache Public License), save one shocking exception: the Dalvik VM, which as the heart and brain of the stack, is proprietary! And we enjoyed some live action too - a demo of Android on the emulator, live! And then Thilak was joined by Hooja and Bhandari for a panel discussion on the various mobile platforms.

The session was highly interactive

with a number of questions from the audience. But what stumped the panellists was the question: "Why are you giving me all this? Is there nothing you can offer that will run on *my* phone?" When the person who had asked the question was, in turn, queried what phone he was using, the reply was: "It's a 1100."

The CXOSummit on Virtualisation

The CXOSummit started with a talk on "Virtualisation and FOSS—What's the Connection?" by Srivathsa N.S., engineering manager, Dell India. He gave some insights into why virtualisation matters for the enterprise. He touched upon the technological aspect and also created the ideal atmosphere for the discussions to follow.

Next up was a panel discussion on "Why Cloud Computing, SaaS and Virtualisation would make business sense, especially coupled with FOSS, at a time when people are looking for cost cuts?" This was chaired by Arun Bhandwal, senior manager of product marketing solutions, Dell. The panellists included Abhijit Das, Shashir Jain, CEO and president, Netcellence Technologies, N.S. Gowri Ganesh, project leader, CDAC Chennai, and P.K. Mishra.

This was a session that had the audience going in full swing, cross-questioning the panellists. The latter rose to the occasion, seeming to enjoy every question thrown at them. Some

The Gurus@OSI

- Ashjit Das, Konulte
- Akash Sinha, engineering student, IIT Madras
- Amit Shah, Red Hat
- Anil Wazir, Microsoft
- Anon Bhardwaj, Dell
- Ashish Mathur, IBM
- B.C. Sekar, HCL
- Bambang Sitohang, Centre for Empowerment of OSS-ITB, Indonesia
- Boudhayan Gupta, FOSS enthusiast & hacker
- Bhuvaneshwar A, CollabNet
- C.N. Krishnan, AU-KBC & NRCFOSS
- C.R. Mathukrishnan, IIT Madras
- D. Jeyakrishnan, IIT Madras
- D. Seetharam, IBM
- Dhruj Sinha, Perot Systems
- Ditya Prakash, Atri Inc.
- Farida Ummari Khan, IIT Bombay
- Gautam Sinha, CTO, Times Business Solutions
- H.R. Mahan, IEEE Computer Society
- J.A. Bhavsar, ITM, Kharghar
- John Smedley, Ingres Corporation
- Kridh Meringkun, National Electronics and Computer Technology, Thailand
- Krish Adlak, TCS
- Lim Kin Chew, SIM University, Singapore
- M. Sankaranarayanan, C-DAC
- Martyn YAE, SIM University, Singapore
- Mathu Ramadoss, Android Rocks
- N.S. Gowri Ganesh, C-DAC
- Narendran Bhandari, Intel
- Naveel Khosla, IBM
- P. Balasubramanian, NIC
- P.K. Mishra, Nixell
- Rajesh Iyer, Essentia
- Rakshit Hooja, IDA Systems
- Radhwan D. Oas, leader, ADomains Project
- S. Ramakrishnan, C-DAC
- S. Srinivasan, NRCFOSS AU-KBC Research Centre
- Sambhaji Na Parthasarathy, Senior Data Centre Consultant, IBM
- Satyashankar Goswami, consultant
- Shishir Jain, CEO, Netpollence Technologies
- Shweeb Hussain, KT Computing
- Sonali Minodha, CTO, OSSCube
- Sridhar Iyer, IIT Bombay
- Srivathsa N.S., Dell
- Steve Lau, Hong Kong OSS Centre
- Subramaniam Bharati & Chennai LUG Team
- Sumeet Arora, Evan Technologies
- Supreet Sethi, consultant
- T. Srinivasan, Cloudiversity
- Thakur Kemer, Wind River
- Tim Clonan, JasperSoft
- Venad Gupta, head, Kean & Ahle
- Venkat Mangudi, Venkat Mangudi Consulting
- Venkatesh Kamaravel, IBM
- Vikash Jha, HSBC Investment
- Vineet Agrawal, co-founder, OSSCube
- Vinak Bhattacharjee, Ingres Corporation

We hope we have been able to include all our speakers in this list. If we have missed someone please be assured that it's certainly not our intention.

We can't 'Thank You' enough!

The organisers would like to take this opportunity to thank all those who played a significant role in shaping OSI.

- Dr S Ramakrishnan, director general, CDAC, for a highly motivating inaugural speech and for encouraging us despite some initial hiccups at the inauguration.
- Dr Rajagopalan, executive director, CDAC, for suggesting Chennai as the venue for OSI and supporting us throughout the planning and execution phases.
- Dr Sankaranarayanan, principal research scientist, CDAC, for helping us shape the agenda for FOSS-enabled education workshop and for giving us guidance from time-to-time on how to improve the activities.
- The entire CDAC Chennai team for not only sponsoring the show, but also assisting us in executing the same. Special thanks to Mr Nagesh Ramaree for supporting us in various aspects of execution.
- The entire IBM team for supporting OSI as Gold Sponsor—an initiative that was spearheaded by Mr Ritesh Segal, marketing manager, Linux and Mr Suresh Wangoo, marketing manager, software group of IBM.
- Dr S Srinivasan, project scientist, AU-KBC, for connecting us with many key people who ended up supporting OSI.
- The entire Chennai LUG for supporting OSI wholeheartedly.
- Mr Ranjit Sengupta, CEO, Advantage Pro, who worked overtime to help us connect with engineering colleges and to ensure that the next generation of engineers were exposed to merits of Open Source.
- Our exhibitors—Linux Academy, Advantage Pro, Learning Linux Centre, Chennai LUG, IIT Mumbai and Open Street Map.
- Last but not the least, all our speakers, who came from all parts of India to share their knowledge, and all of you who took out time to attend OSI.

asked about the connection of FOSS with cloud computing, whereas others raised questions about the security of the data. Cost was another factor that was discussed. However, there seems to be some misunderstanding about the 'free as in freedom' aspect. Some folks from a leading India ISP reported about a perception that where things are 'free', there's automatically a lot of concern about security. While the panel covered all worries about security, they stayed away from going into the actual implications of what 'free' meant in FOSS. There is obviously still a lot of confusion about this magic word even among open source users.

Talking about the importance of staying safe, one panellist's memorable quote was: "It's good to be passionate about something, but it's even better to be bullet-proof." Implying that no matter what technology you use, security is something you should keep as top priority.

The third session of the day discussed the various FOSS virtualisation technologies in use. The panellists included Amit Shah, KVM developer at Red Hat, Shweeb Hussain, entrepreneur and free software developer, P.K. Mishra and Venkat Mangudi. The discussion was moderated by Gautam Sinha, CTO, Times Business Solutions.

When a member of the audience raised the important request for help in choosing one from among the various

technologies, Shah covered all bases. He said that it was most important to always ask for a trial, to test two or three technologies, and see which one meets your needs the best. He warned against blindly going after only one solution without testing others.

This session was apparently 'hijacked' by Novell's P.K. Mishra, making it one of the most entertaining at the event. Given that the audience was shooting off question after question, Sinha had to extend the session well beyond its time limit.

Dedicated to education

Targeted at the academia, 'FOSS-enabled Education' was the third parallel track of the day. Discussions were held on topics like 'A FOSS-powered IT set up', 'FOSS and education', 'Teaching FOSS', 'Open source in higher education' and so on. Prof. Bhavsar, Group Head (IT), Institute of Technology and Management, Kharghar (a place in the suburbs of Mumbai), said that while paid software has its own advantages, very few educational institutions can afford to buy and use paid software to meet all their needs. This is best met by open source software, which has its own merits. Being free is not its only attraction—also inherent in the open source philosophy is the freedom of a distributed community of programmers to modify and improve the code. This model exemplifies the best of humanitarian efforts for a social cause.

No gatherings
or get-togethers,
yet 6 billion
conversations
everyday.

Making the impossible possible, is a way of life at Yahoo!. Small or big, whatever be the idea, we bring it to life. No wonder, Yahoo! has constantly conceived and delivered worldclass products that touch the lives of over 500 million users. Whether it's email, news & entertainment, sports, music or just about anything else on the net, people turn to Yahoo!. Making it the world's most visited site and among the most well-known brands globally, Yahoo! India R&D is a strategic Center that contributes in making Yahoo! what it is.

We are presently looking for technology professionals for the following positions:

Tech Lead - Virtualization R&D

The Cloud Computing group at Yahoo! is chartered with making our data center environment friendly and more efficient. We are working at cutting edge of technology to create Internet-scale solutions that meet these goals by leveraging techniques like virtualization. As a member of this team, you will be part of the next wave of innovation at Yahoo! to transform our data center into Cloud Service Delivery Platform.

The ideal candidate should • Be an M.Tech/MS Computer Science (or equivalent) with minimum of 4+ years industry experience • Be proficient in designing complex enterprise class/Internet scale applications addressing requirements like high performance, maintainability & service ability • Have good understanding and hands-on experience of design patterns and best practices for large scale software development • Be expert in Java programming language on JVM/Linux platform • Have conceptual level understanding of virtualization technologies • Be someone who can bring in fresh ideas and doesn't see the world the way everyone does

Principal Engineer - Performance Engineering

The Performance Engineering team is an expert group that studies, analyzes and improves performance and scalability of key Yahoo! applications. This includes Yahoo! Search, Advertiser-Publisher Technology and our Cloud Computing Infrastructure - Apache Hadoop. This group co-works with product development teams to address top performance & scale challenges ranging from crunching peta-bytes of data every-day to serving relevant search results in mill-seconds. Yahoo! is currently building presence for this expert group in Bangalore and is looking for bright senior engineers. The ideal candidate should possess skills to dive deep into design, code, operating system, and hardware-level details of software execution.

The ideal candidate should • Be an M.Tech/MS Computer Science (or equivalent), 3+ years industry

experience • Have proficiency in C/C++ or Java multi-threaded programming for Java programmers, knowledge JVM internals is a plus • Have solid fundamentals in Operating Systems, Networks and File Systems. Knowledge of Linux Internals is a strong plus • Have good understanding of modern multi-core systems, cache and memory architectures • Have strong distributed computing concepts, as applied to large scale (1000+ nodes) systems experience with Apache Hadoop or Map-Reduce is a plus • Background in Performance Optimization or High Performance Computing (HPC) is a plus

Tech Lead - Linux Engineering

The Linux team is the "Linux Center of Excellence" for all of Yahoo!. This team is responsible for Linux strategy, support and development for virtualized and non-virtualized environments. This team works very closely with Yahoo! properties, Linux vendors and the open source community to provide a stable, secure and scalable infrastructure to build complex Yahoo! applications upon and comprises of a group of highly motivated and energetic kernel systems and release engineers spread across US and India. As a Linux engineering lead, you will contribute individually to Linux and Linux Kernel adoption in Yahoo! by participating in various activities including • Migrating applications from FreeBSD to Linux • Providing Linux kernel support • Packaging popularly used public domain Linux software for end-use in the Yahoo! environment

The ideal candidate should • Be a B.Tech Computer Science (or equivalent), 7+ years industry experience, M.Tech/MS is a plus • Have extensive experience developing software on a Linux environment, Unix / Linux kernel development experience is a strong plus • Have a good understanding of packaging and releasing software on Linux, including RPMs • Have excellent software development, debugging and testing skills in C/C++ • Background in porting source code to Linux is a plus • Scripting skills in Perl/Python & Unix shell programming is a plus

Mail in your resume to careers@yahoo-inc.com
For more information log on to <http://india.yahoo.com>

YAHOO!
INDIA
RESEARCH & DEVELOPMENT

Sridhar Iyer, associate professor, Dept. of Computer Science, IIT Bombay, talked about two case studies of a FOSS set up in the institution, Project Oscar and Project Ekalavya. These projects opened the students' eyes to the benefits of FOSS and also trained them in developing software in open source, he said.

"Flexibility is the key benefit of FOSS apart from the low cost associated with it. Also, open source facilitates creativity in the students, since they can directly put their ideas into practice as there is no upfront cost, proprietary rights and copyrights associated in the initial set up," said D. Janakiram, Department of Computer Science, IIT Madras. And in the educational scenario, open source provides an immense opportunity to develop software, right from designing it to offering support.

The panel discussion on '*Adopting FOSS in education—The road ahead*' was chaired by Dr M Sasikumar, associate director (research), CDAC, Mumbai. He said that compared to other industries, only in the IT industry was the 'raw material' vendor specific. This creates problems for academia, since they cannot concentrate on one area. He also said that industry and academia should join hands together to promote open source and to solve the problems associated with it. C.N. Krishnan, program director, MIT Campus, Anna University, said that in the present education scenario 85 per cent of the curriculum concentrates on theories, and only 15 per cent on practical aspects, while open source is all about 'practicalities'. So there should be a revamp in the curriculum.

FOSS For Everyone

The final day of the seminar was dedicated to the idea of 'FOSS for Everyone'. Boudhayan Gupta, a fourteen-year-old FOSS fanatic, spoke on how he fell in love with FOSS. He said that FOSS actually makes more sense because we live in the age of Liberty, Equality, Freedom and Fraternity, and FOSS gives us all these rights. He also said that FOSS offers the freedom to do whatever you want to with the software, such as take it apart, hack it, and all without being restricted in any way either legally or financially.

Gautam Sinha, CTO, Times Business Solutions, spoke on why TBS chose FOSS over proprietary solutions. He said that FOSS itself is a powerful word, which has two strong words in it, 'free' and 'open'. "We were working with proprietary software and we didn't have too many issues with it. But when we moved to FOSS we had advantages over cost and our productivity also increased," he said. He added that the heavy traffic of TimesJobs.Com, which has a database of ten million users, and SimplyMarry.com is smoothly managed by FOSS software, with sub-five second response times for each query.

The Chennai LUG initiated a panel discussion on the topic '*Free as in Freedom Distros*'. Bharati, a member of the Chennai LUG, said that if you choose proprietary software, you have only different versions of one product like Windows 98, Windows XP, Windows Vista and so on, but if you opt for open source, you have a variety of choices like Ubuntu, Fedora, Gentoo, Mandriva, BOSS and so on. The panellists presented the case of each of these distros, one by one.

Krish Ashok, head, TCS Innovations Lab- Web 2.0, shared a presentation outlining the Open Source movement that's slowly gaining momentum within TCS. Apparently, they plan to launch a corporate blog very soon, which is going to be based on Joomla. And as we all know now, the brand new website of the Tata Nano has also been developed on the same CMS.

The event also witnessed the

Another team mate of ours, LFY's official Tweeter, is obsessed with his newly acquired BlackBerry. While he was live tweeting from his phone, he had no idea about his surroundings. Once, he was too busy to get his hand baggage tagged by the security and had to rush back, at the very last minute, while the whole plane was waiting for him before the take-off.

entertaining and thought provoking talk by red1 (that's Redhuan D. On, the chief developer of the ADempiere ERP Suite) who jokes that he is the next Linus Torvalds-cum-RMS from Asia. He criticised developers who try to keep control of their projects by stealing the bug reports filed by others. To quote this outspoken gentleman, "If I submit a bug fix, he will say, 'Oh I saw it, it's not good.' And then he will peep into my code, rewrite it and publish it as his own code, stating now that it's perfect." He also said that you need not worry about some open source project being made closed source in the future, because once the code has been GPLed, you cannot close it. Of course, subsequent versions can have a closed licence, but you can always fork the last GPL version and continue development as a separate project yourself.

"So may the fork be with you!" That was his 'call to arms' during dinner, too!

Konclusion

The event, by our standards, truly rocked as we succeeded in bringing together different audiences and we believe everyone benefited from all the interactions. The time schedules were tight, but we managed to pack in a whole lot, in the short period of time. What we've been able to bring to you is just some of what happened. You needed to be *right there* to soak in the infectious energy and buzz around open source. *That's* not transferable to newsprint! 

Authors

Swapnil Bhartiya, K T P Radhika and Boudhayan Gupta. Sayantan Pal, Jesus Milton Frousseseu and Atanu Datta contributed to the article: Free As In Beer, Too! Only

On our last day in Chennai, right after the conclusion of OGI'08, I met a college student who had no idea about open source before he'd attended the event. We showed us a journal that he had filled up with actual notes of the whole event. He planned to create a presentation for his college mates who'd missed OGI. We feel, we could have given you a better account of the sessions than we have managed to, in this issue. -P

Create Slide Shows with 3D and More

Have you ever wondered how to create a stunning slide show with audio, 3D effects, animation, mind-boggling slide transition effects and more? And you must have also wanted a simple way to create it. Then it's time to give Manslide a try.

Manslide is a Qt4 application that can help you create a 3D video slide show and afterwards save the creation as MPEG2, XVID, FLV or DV formats to either run on a video player application on your computer, or upload on YouTube. Yes, of course, you do have a say about the video quality and can adjust it to suit your requirements.

Before you get started, you need to take care of the requirements first:

- QT4 >= 4.3
- smilutils >= 0.3.2
- netpbm >= 10.34
- mjpegtools >= 1.9
- sox >= 13.0
- ImageMagick
- vorbis-tools >= 1.1.1
- mplayer / mencoder

And, of course, keep your photographs and a pleasant voice (or music) track ready, if you decide to include audio in your masterpiece.

Getting Manslide

Manslide 2.0.3 is included by default in Mandriva Spring 2008 DVD. You can download the latest package from linux.softpedia.com/get/Multimedia/Graphics/Manslide-18081.shtml. There is also a standalone package on the site. Uncompress the application, change into that directory and run `./Manslide`.

Let's get started

Step 1: The default language in Manslide is French. In the first window that opens, change your language to English from the drop down language setting field. You will be prompted to enter the name of a folder to save some Manslide animations and slide transition effects.

The top black menu bar has buttons titled *Project*, *Pictures*, *Parameters*, *Soundtrack* and *Export*—and this is essentially the order that has to be followed in your slide creation.

Step 2: Restart the application. Under the *Project* tab, click on *Start a new project*. Now enter the name of the destination folder in which you want to save your final creation (Figure 1).

Step 3: Now click on the *Pictures* tab and then on the *Add* button (the one with a PLUS sign on the left side of the window). Navigate to the folder, where your images are placed. Now,



Figure 1: Choose your destination folder to save your creation



Figure 3: Image effects



Figure 2: Adding slides



Figure 4: Slide parameters

click on an image and then on the *Insert slides* button (which is present on the middle black bar on the right hand side). The added slides will be shown at the bottom. When done, click on the *Close* button. Refer to Figure 2 for a better idea.

Step 4: Now is the time to add some groovy effects to your slides. Click on the *Effects* button (near the PLUS button) and choose from the various effects like flipping, horizontal/vertical mirror, emboss, fade, etc. Refer to Figures 3 and 4.

Choose the starting slide as 1, and the last slide's number should be equal to the number of slides you have chosen. By default, the slide viewing period is 3 seconds and the transition takes over 1 second. Change these settings if you want to. Now go to the transition types, and you'll find the drop down menu is mind blowing! Selecting from the over 300 transition types available is no easy task.

If all this isn't enough for you, you can also add colour backgrounds and background images to your slides by clicking the respective buttons.

Step 5: Click on the *Parameters* tab. You will




Figure 5: Presentation tab: Effects and more slide effects

be presented with more slide and animation options like motion blur effects, superimpose animation on a slide, fog effects, add title to a slide and slide presentation (new slide entry effects like window, vertical/horizontal double plane, animated flag, cube, etc). Refer to Figure 5.



Figure 6: Adding a soundtrack

Step 6: How about adding a soundtrack? This could be a narration you've recorded, or just a music file. Click on the *Soundtrack* tab and then on the *Import* button (Figure 6). Choose your track from the navigation window that opens. There are even sound effect options like fading in and fading out, and volume control. Further, even though the length of the soundtrack could be more than the slide show's length, Manslide will adjust the track length to the slide show's length, automatically.

Step 7: And finally... I can't wait any longer for my own slide show video! Click on the *Export* button. Enter the name of your project in the slide show title field and the type of codec (MPEG2, XVID, DV or Flash) from the drop down menu. You can also select the video quality (high, standard or low) for your project. Refer to Figure 7.

When all is done, click on the *Start* button and find your masterpiece in the destination folder that you selected at the start of your project. The creation time depends on the number of slides and video quality chosen.

Note: On the vertical pane on the right hand side of Manslide are controls to orient your slides in X, Y or Z direction. Also available are two flags to start and stop animations. To use the effect of these flags in your slide, click on one or more of the X, Y or Z buttons, and then click on the *Start* flag. Right




Figure 7: The Export button

click on your slide and hold the mouse button down to change the orientation of the slide. Once done, click on the *Stop* flag.

If at any time you wish to stop your unfinished work, click on *Project* → *Save Project*. The saved file will have extension *.msz*. To continue sometime later, run Manslide again and navigate to *Project* → *Open existing project*.

Caution: Manslide is not verbose as far as reporting missing dependencies is concerned. On installing Manslide from the Mandriva Spring DVD, I found that *netpbm* and *smilutils* were not installed, which resulted in the final video slide show not being created. However, on installing the required dependencies, all problems were solved.

The final touch

I only have one word to define this masterpiece of an application; Fantastic! Time to call over my friends who're still stuck to their Windows for some tea and cookies (my grandma's masterpiece), with the secret intention of playing them my masterpiece video. After all, it's Windows' envy, Linux's pride. 

By: Nelson Lobo

The author is a lecturer at St Xavier's College, Mapusa, Goa. He can be reached at lobonano@gmail.com



is now



Open Source. Open Standards. Open Minds

BITS Pilani



present

apogee'09

23-27 march

XXVII International Technology Festival

We Thank all our Sponsors
for being a part of Apogee'09

apogee'09

Titan Sponsors



Gold Sponsors



Silver Sponsors



Platinum Sponsors



Exhibition Partners



Partners



Food & Beverage Partners





OPEN SOURCE INDIA CHENNAI



FOR YOU & ME

- 18 Dell Has Everything to Beat Apple
- 20 Debian Lenny: Brewed to Perfection
- 24 Solaris for your Desktop
- 28 Open Source India '09
—An Event Report
- 36 Create Slide Shows with 3D
and More
- 40 "Because of the much lower cost,
open source should do very well in
this economy."
—An interview with Bob Young,
co-founder of RedHat
- 42 FOSSConf'09: Three Days of
Software Nirvana!
- 45 The Delhi Drupal Meetup v1.0
- 46 The Open Sesame of Business
- 48 Teeworlds: Ready to Jump the Gun?
- 50 ICFOSS: Wide Open Windows
- 52 Why Not Carry A Pen Drive
Instead?
- 94 Getting Your Hands Dirty With Deb
Files
- 99 Envisage'09: Festival To Celebrate
Big Ideas



DEVELOPERS

- 72 USB 3.0: Buckle Up for Super Speed
- 74 Building IM Toys
- 76 'Good' Coding Practices
- 80 The Intricacies of the System V
Semaphore
- 84 A Walk Through Subversion 1.6
- 98 Book Review: Embedded Linux
Primer



“Because of the much lower cost, open source should do very well in this economy.”



Bob Young has played a key role in building Linux into a household name. In 1993, with Mark Ewing, he co-founded Red Hat and was the CEO of the company for several years. After leaving Red Hat, he started a print-on-demand website called Lulu, where content creators can sell their books, comics, movies, or any other content that can be digitised and sold over the Web. Young also owns the Hamilton Tiger-Cats, a Canadian football league franchise that he purchased in 2003. In this interview with LFY, he talks about Lulu.

Q What attracted you to open source and eventually to co-founding Red Hat?
I was unemployed.

Q How did you arrive at a revenue model in open source? Did you have a role model?
Cygnus Solutions was doing reasonably well selling custom engineering services that adapted open source languages and compilers for new processor platforms, mostly for the embedded market.

And there were a few very small Linux competitors, such as Yggdrasil, Turbo Linux and SuSE, who were experimenting with business models at the same time Red Hat was starting up. Only SuSE has survived today as a part of Novell Corporation.

Red Hat's innovation was to

understand that users of proprietary software wanted control over the software they were using, but the suppliers (Sun Microsystems, IBM and Microsoft, among others) were not prepared to give them control. Which gave Red Hat its business opportunity. As [Richard] Stallman would have said, "It was not about price, it was about freedom."

Q Since you mentioned RMS, how do you differentiate between free software and open source software?

I don't! You'll have to ask Richard Stallman and Eric Raymond. They have had many public debates on the fine points of difference. I see open source as simply a better marketing term than 'free' software.

Q Today, more and more people and organisations are migrating to Linux. How do you view this trend?

Linus Torvalds used to joke when asked about the future of Linux—he would say, “World domination!” Seriously, in an interconnected world made possible by the Internet, the reliability of the systems that are connected together becomes ever more important. Only when the managers of those systems have control over their systems (that is, source code and a licence to modify it) can they build proper reliable networks. Open source is simply a better model of software development and deployment for the Internet age than the traditional proprietary model.

Q From an open source software platform, how did you enter the digital market and publishing? What is the story behind setting up Lulu [lulu.com]?

Most open source software is built on a sort of barter system, where professional engineering teams trade software that they develop for a greater amount of software that the broader community develops. On the other hand the fact is there are many unpaid volunteers. That these valuable contributors were unpaid used to worry me. Similarly, I ran into many capable authors who could not earn money from their books as they had been refused by publishers. Using the Internet, Lulu can help authors sell their work to their readers for trivial cost.

Lulu was founded as an effort to enable creators to get paid for the work they do, if they need to get paid. In particular, documentation writers need a business model if they are to continue to write up-to-date revised documentation for all the open source software this worldwide community of developers is producing. We want to help readers identify the long tail products.

Q So, there is some connection between Lulu and open source...

We like to think that we are enabling a form of ‘open publishing’ where the authors have control of the publishing process and do not have to ask permission of a publisher to make their work available to the market. Much like open source enabling the engineers who use open source software to do what they need to the software, without asking permission of the proprietary software company.

Q What is Lulu’s vision?

As Linus would say, “World Domination!” :-)

Q Tell us something about Hamilton Tiger Cats. How do you connect open source, publishing and football?

Not much! Although there is an interesting connection in that all three are very much “community” oriented projects. Tiger-Cat fans feel strongly that it is ‘their’ team as it connects them to each other and to the community they live in. In the same way, Red Hat Linux users feel that they own Linux not Red Hat. And they are right in both cases.


Q In these times of an economic meltdown, how do you think users respond to open source?

Because of the much lower cost, open source should do very well in this economy.

Q Is the economic upheaval affecting the Linux market too?

Not as badly as the proprietary software vendors :-)

Q What challenges do you see for people migrating to Linux?

None! Well, okay, it is a new technology. As with any new technology, there is a learning curve. 

By: KTP Radhika

The author is a Linux admirer and correspondent at LFX's Bangalore Bureau.

poor JOB MARKET?
Campus Placement?

Future is OPEN

Now, time to Upgrade your Skills. Register...



RHCE

Free Training for RHCE Exam
RHCE Exam - 14 & 30 Apr 2009

First time in INDIA
LPI Certification
Available now ...



- LPIC 1 & LPIC 2
- UCE (Ubuntu
Certified Engineer)

Courses Starts from Apr 09
Register Now ...



Zend Certification

ZCE - International Certification for PHP
from Zend Technologies - USA

Save
Up to **40%** Now

Sp! Offer Register Zend GET
LAMP Pro Training (40 Hrs) FREE
Valid upto 24 Apr 2009



Future is open

Lynus Academy Pvt Ltd

Chennai - 044-25021630,
42171278, 9952966527

Pondicherry - 0413-2276233,
4202121, 9952966528

contactus@lynusacademy.com
www.lynusacademy.com

* Conditions Apply



FOSSConf'09

Three Days of Software Nirvana!

FOSSConf'09 brought together people who shared one common interest—free and open source software. Apart from interesting talks and demo stalls, it was a chance to network among those with a passion for freedom.

*T*hiagarajar College of Engineering (TCE), Madurai, hosted a three-day national conference on free and open source software, dubbed FOSSConf'09, from February 27 to March 1, 2009. Prof S.V. Raghavan, IIT, Madras, inaugurated the conference and highlighted the role of FOSS in the various research initiatives across the world.

FOSSConf'09 was jointly organised by TCE, ILUGC and NRCFOSS. FOSS enthusiasts from across the nation who

contributed or who are interested in contributing, came to the event to meet, discuss, brainstorm, innovate and build on existing and new FOSS projects.

Motivated by the philosophy behind free software [www.gnu.org/philosophy/free-sw.html], the GNU/Linux User Group of TCE (GLUGOT) was started in December 2003 with the objective of promoting free software and the spirit of cooperation among the students. The group contributes to the community in the form of awareness meetings and

technical sessions. It regularly conducts training programs to teach students about free software-related technologies. TCENet is the FOSS-based ERP package developed in-house to automate the day-to-day activities of the institution [tccnet.tce.edu].

In addition, there are many on-going interdisciplinary research programmes in FOSS. GLUGOT has been organising a Software Freedom Day (FStival) since 2004 to educate college and school students and the general public, in and around Madurai. This year, FStival is scheduled to be on September 12, 2009, where demo stalls and technology sessions on various FOSS tools will be conducted [fstival.tce.edu].

This group also facilitates the meeting of GLUG-Madurai at the TCE Campus. The GLUGOT members have developed software projects like TuxEdo—a distributed chess engine, Instant Messenger Jammer, and GetEz, which is an interface to the wget tool for GNOME. These projects are available on the Web as free downloads. FOSSConf09 is the first National Level Conference organised by this group.

FOSSConf'09 spread awareness about free software among new users and contributors. It was an opportunity for students of various age groups and domains, professionals (from both IT and non-IT fields) and the curious, to get introduced to FOSS, interact with each other and share their knowledge. The conference also showcased projects, research presentations on FOSS and FOSS-based contests emphasising the importance of free software technologies. There was Wi-Fi blanketing the venue, and high-speed connectivity was available to all participants at the event. In fact,

people were encouraged to bring their laptops, or Wi-Fi-enabled devices to better experience FOSSConf'09.

Over 10 different organisations hosted several projects developed in FOSS. The project stalls hosted a number of technical projects. These included TCENet (ERP for educational institutions), TAMS (a biometric and RFID-based attendance monitoring system), TOSS (a customised distro), TCE High Performance clusters, NLP Tools, Snort, an IDS, an automated Internet bandwidth controller for Linux, an SMS travel guide, a virtual music generator, a portable music player, Wimax implementation using estimation algorithms, performance analysis on Web traffic, an online vehicle insurance payment application and so on.

There were 20 demos covering the basics of Free/Open source software, educational tools, Linux thin clients, a mobile terminal, Django, SOLR Enterprise Search Server, Drupal, security power tools in Linux, auto installation of the Globus tool kit, Yavarkumana Menporutkal in Tamil, Linux commands, Office tools, live CDs, GCC, the GIMP, bioinformatics, Wireshark and other packet sniffers, GNOME, KDE, virtualisation using Sun xVM VirtualBox, etc.

In addition to the project and demo stalls, various talks were arranged in three parallel tracks. These focused on the nuts and bolts of how people can participate in the community. According to the speakers, there are plenty of reasons to be a part of the process. Helping a free software project can be a way to learn skills, explore ideas and their implementation, have fun, create employment opportunities, and work for the



India's Premier Linux and Open Source Software
Training Academy, Looking for Franchisee ...
JOIN Us - future@lynusacademy.com

Future is
OPEN
www.lynusacademy.com



▲ Prof S.V. Raghavan inaugurates the FOSSConf'09



▲ Principal Dr V Athrikumar gives the inaugural address



▲ Demo stalls



▲ Volunteers prepare charts for display



▲ Valedictory function at the K.S. Auditorium, TCE



▲ Participants at demo stalls

common social good.

Writing code is the first and foremost way of participating, and speakers dispensed a fair amount of advice on how that is best done. But they also took time to point out the many other ways to help, most of which do not require programming skills. These range from reporting bugs, writing documentation, translations and localisation, creating artworks, and helping to maintain the infrastructure needed by free software projects.

Localisation was highlighted as an area in constant need of work. India has a long list of languages to translate into, and Indians are the only ones who are well positioned to get that work done.

Day 1

The first day covered topics like freeing geographical information, WLANs and open source software, virtualisation, bioinformatics (from the point of view of open source), kernel module programming, Real-time Linux, open text books, open desktops, BioPuppy (a minimal Linux application for bioinformatics) and GlassFish.

Day 2

On the second day, those who attended the FOSSconf got a chance to listen to the talks on Ruby on Rails, open source in data integration, how FOSS could help you parallelise your programs, an introduction to the RHCE Linux Web, the Apache set-up and configuration, open electronics labs, *Badam Halwa* of Embedded Systems, C under Linux, job opportunities in open source, LTSP (the Linux thin client), building and using a custom C library, Magneto (an e-commerce CMS), how to create and implement a patch, etc.

Day 3

The third day featured sessions on Perl for beginners, GNU/

Linux network tools, network hacking the open source way, Eclipse—a programmer's weapon, software requirements and design specifications, Valgrind, free software, shell scripting, open source vs open systems, etc. The valedictory function was presided over by the principal, TCE, who distributed prizes for the Best Speaker, the Best Project and the Best Demo Stall.

The official website of the conference was fossconf.in and as soon as registration opened on it, there was an overwhelming response to register for talks, projects and demo stalls. All the 30 selected projects were developed in FOSS, published on the website and were ready for distribution too. The slides of the 35 speakers have been made available on the website.

No entry fee was collected for the conference. FOSSConf'09 was organised and conducted entirely by FOSS community volunteers. The event was completely non-commercial in nature. Around 1,500 visitors attended the conference over the three days. The organisers had taken the utmost care to make this conference highly beneficial to the participants and a perfect place for knowledge sharing.

Speakers, delegates and visitors benefited from this cooperative development model and hence proved FOSS, otherwise called 'Society responsible software', a successful methodology and a part of a movement. For more details visit glugot.tce.edu or join the mailing list, which is the best place to ask about things related to the event lists.tce.edu/mailman/listinfo/glugot.



By: Dr S. Mercy Shalinie

The author heads the Department of Computer Science and Engineering, Thiagarajar College of Engineering, Madurai. She can be reached at shalinie@tce.edu

The Delhi Drupal Meetup v1.0

This meetup saw Drupal enthusiasts turning up from different parts of the country, such as Srinagar, Meerut, Bijnor and the NCR region.

On March 1, 2009, LUG@IITD hosted the Delhi Drupal Meetup at IIT Delhi. It was the first of its kind in the country, with the aim of ensuring contribution from all the participants. The meetup saw a turnout of 70 people with laptops, all set to code for Drupal. Anyone attending the meetup had to write a Drupal theme or module. This interesting event was managed by just four volunteers.

Kicking off with Drupal theming

The event was kick-started by Vivek Khurana's talk on Drupal theming. He walked the audience through how a Drupal theme is structured, the core components of a Drupal theme, the concepts of blocks and regions, followed by templating engines available with Drupal themes. The talk had examples of a theme with a template engine and one without. The talk concluded with how to create separate designs for specific pages and creating a theme from scratch.

Initialising the module

Pratul Kalia spoke on Drupal module development. He explained the Drupal module system, its architecture and the Drupal menu system. He used examples to explain things like the Drupal module node API, using external PHP modules in Drupal modules, and how to start creating a module from scratch.

Clash of the Titans

Subsequently, participants were given a break to think about the ideas discussed, over a cup of tea and *samosas*. Participants interacted with each other to prepare for the real coding. During the post tea session, all participants sat down for a contest on theming and module development. While the code competition was going on, Supreet Sethi and Piyush Verma took the lead in guiding the Drupal newbies in development problems that they were facing.

While basic support was provided by Supreet and Piyush, Vivek, Mir Nazim and Pratul helped participants with the more advanced issues related to database design and debugging PHP code.

"And the winners are..."

After two hours of coding, a couple of participants were ready with themes, but the module developers demanded extra time. It was unanimously decided that module developers should be given two days to submit their entries.

The theme contest was judged by the participants that very day, and Kaushal Kishore won the prize for the best theme. He won a bag full of goodies and a one-year *LFY* subscription.

The module development contest was judged three days later and Ravi Gupta won the first prize. He developed a module to find the difference in the database schema of two Drupal installations, developed from the point of view of easily moving features from the development server to the production server. He, too, won a bag of goodies and a one-year *LFY* subscription. The judges decided to award a consolation prize to Amit Sethi, a student of IP University, for his module that allows linking keywords in a Drupal node to Wikipedia and Google search results.

The bags of goodies were sponsored by OSSCube, while Srijan Technologies sponsored the snacks for participants, and dinner for the volunteers and speakers. *LFY* magazine sponsored the one-year subscription for the winners. The event was organised by Mayank Kumar, Gajender Khanna, Sangeet Kumar and Vivek Khurana of LUG@IITD. 

By: **Vivek Khurana** reporting on behalf of LUG@IITD.

• Amit Sethi with his prize





The Open Sesame of Business

The recently held seminar on open source at the Bangalore International Exhibition Centre was an attempt at creating awareness on using open source for business enterprises. The seminar covered the challenges companies face while making an 'open' shift, and the ways to tackle them. It was jointly organised by the Confederation of Indian Industry (CII) and the Ministry of Micro, Small and Medium Enterprises, government of India.

Lack of awareness prevents enterprises from enjoying the benefits of open source applications. This was one of the conclusions of a seminar on open source software and SMEs, held in the last week of February in Bangalore. "Open source is a great platform for industries. For instance, even the New York Stock Exchange runs on open source applications. Micro, Small and Medium Enterprises (MSMEs) can follow open source as a developmental model as well as a business model," said Nandu Pradhan, president and managing director, Red Hat India, while delivering the keynote

address at the seminar.

According to Pradhan, companies like Red Hat take the free resources from the open source community; develop, integrate and test them; and then make them ready for enterprises. "We provide the support. Unlike proprietary software, which sells the product, open source products are free, and services are charged," he said.

In his special address, M. Sasikumar, associate director (research), C-DAC, stressed on the freedom that open source software delivers. He said that lack of awareness was the main barrier to SMEs opting for FOSS. "In large scale sectors,

open source is gaining momentum, but SMEs have not been targeted yet. Open source will benefit SMEs for multiple reasons: the first is cost—the initial cost for software is negligible, compared to proprietary software. OSS provides free, good-quality technology solutions. And since it is open, users can potentially adapt it to their needs and customise it according to their requirements. Providing support to SMEs in open source is itself an opportunity for a start-up,” Sasikumar said.

However, Vipin Singh, director—EDCS, e-governance department of Karnataka had a different take. “More than open source, SMEs should first start using information technology. IT penetration is very low among SMEs. Open source brings in huge opportunities. It may not solve all the problems. But it is very secure and bugs are removed more frequently than in proprietary software,” he said, delivering the inaugural address. Vinay L. Deshpande, chairman and CEO, Encore Software and Surendra Kumar, chief executive, Precision International also spoke during the event.

The inaugural session was followed by a round-table session on open source. R K V S Raman, senior scientist, CDAC-Bangalore, chaired the session, titled ‘OSS awareness, access, application and acceptability in MSMEs’. Manish Malhotra, director and country-head, software, Sun Microsystems India, Venkat Mangudi, founder and CEO, Venkat Mangudi Consulting, and Arun Mehta, OSS consultant, also spoke at the gathering.

Raman said that open source had several clear advantages. It has no licence fees, and no vendor lock-in; it is almost entirely based on open standards, while being available in various sizes. It has great flexibility for the various needs of a business, and users will get support from small firms or

even larger companies like Red Hat or Novell.

Malhotra pointed out that a product need not be considered an open source product if any one component of it is open. Open source only means that the source is available and code is distributed freely. “The market is shifting to open source, and many companies are using open source for mission-critical applications. Ubuntu is emerging as one of the most popular Linux desktops,” he noted. Vendors can make money by providing subscriptions, support, services, advanced training and so on. “Total cost of ownership, hardware costs, cost of software, licence costs, annual security costs, etc, can be reduced or brought to nil by using open source,” Mangudi pointed out.

“Today, writing software is like walking over a mine due to ‘intellectual property’ issues. As a small company, you also have to protect your own intellectual property. If one is working on OSS, one should not have to worry about ‘intellectual property’,” Mehta said. FOSS has several advantages like inherent UNIX security, direct interaction with developers, cheaper hardware requirements, easy installation and upgrade process, excellent Web facility, and cross-platform support.

This session was followed by a panel discussion on *OSS-Technology: Convenience vs Business Risk*, chaired by Sasikumar. Hemant Verma, deputy director, CII-SME Division, Malhotra, Mehta, Raman and Mangudi participated in the discussion. The panel discussed the challenges of evaluating and adopting IT as well as open source software for business requirements.



By: **KTP Radhika**

The author is a Linux admirer and correspondent at LPF's Bangalore Bureau.

redhat®

TRAINING PARTNER

RHCE / RHCSS Exam Centre

**Hundreds of hardcore
Industry Professionals
Vouch for it.**



**RHCE only @
ADVANTAGE PRO**

- Pioneer in LINUX training since 1999.
- Chennai's #1 RHCE training Centre
- 100% results every month.
- Many centum students.
- Also get trained on My SQL - CMDBA, CMDEV, PHP, Perl, Ruby, Python, Ajax...

**Next RHCE/RHCSS Exam on
13th & 27th April'09**

Don't wait.

Register Now.

Let the success vouch for you!



ADVANTAGE PRO
IT TRAINING DIVISION OF
Vectra Technosoft Pvt. Ltd.
Open Source Academy

**Regd. Off: Wing 1 & 2, IV Floor,
Jhaver Plaza, 1A, N.H. Road,
Nungambakkam, Chennai - 34.**

Ph : 28263529 / 3530 / 3540

Telefax : 28263527

**E-mail : enquiry@vectratech.in
www.vectratech.in**

TEEWORLDS

JUMPING THE GUN



Ready to Jump the Gun?

If you fancy 2D multi-player shooter games, check out Teeworlds!

Tux boxes were believed to be best for programming, hacking and learning, but were never considered good for gaming. Yet, there are a few users who love to play good games on their Linux systems. So if you are a gamer or like to play games once in a while, plunge in to this review. Here's a game I found while searching through my distro repositories.

If you care to take a look at Figure 1, that, my friends, is Teeworlds (earlier 'teewars')—a 2D retro multi-player only-shooter game.

First impressions

A look at the screenshots at the game's home website made me feel this could be some adventure genre, role-playing 2D

side scroll game. I thought it would be awfully slow and I sure didn't want another SuperTux on my machine.

A little more Googling revealed that I was wrong; I am glad I was.

Game play

Totally unexpected things showed up when I first logged in to play—weapons from games like Quake with a toonish style, grenade launchers, pistols, and multi-fire laser rifles!

If that were not enough, some maps even feature a katana; reminds me of *Ninja Gaiden*. Wonder who carries these weapons? Tiny, cute little creatures you wouldn't want to kill. Don't expect everyone to feel the same way, though.

On Internet servers, if you are new,



Figure 1: One-on-one death match between Krish and Shona

you'd be killed even before you blinked an eye. In case you haven't noticed, the tag line of the game reads 'Jumping the Gun'.

The default inventory is a pistol and wooden mallet. Other collectible weapons can be reloaded by walking over them. The same goes for life and shield. While the mallet proves very effective during close combat, it is your instincts and reflexes that save you if your opponent has picked up the katana or laser rifle. The game is pretty much in instagib mode unless the weapons are lost.

'Hooks' and 'double-jump' help fast manoeuvres and in dodging attacks.

The game allows chat (useful to taunt your opponents) and an 'emoticon' system to convey your mood to other players. The current supported game modes are *Death Match*, *Capture the Flag* and *Team Death Match*. You can also choose to be a spectator and for people who love one-on-one gaming, you can modify the TDM to allow only two active players.

A little practice on my LAN server and then a few matches on Internet servers were enough to make me addicted to Teeworlds.

Interface

The game menu sports a nice tabbed interface to choose between the Internet and LAN servers. The settings allow you to choose your player skin including options for custom colours; the controls feel like a standard FPS game; the graphics options enable choosing from different resolutions – like the full screen or disabling high resolution details if your graphics card isn't for the game.

There is no joystick support yet, and unfortunately, it is not on the developers' priority list either. Thankfully, re-mapping of keys is allowed.

The sound is satisfactory. Some improvements could



Figure 2: Survival of the swiftest




Figure 3: Snowy grave

include ensuring the pistol and rifle do not sound so similar, and adding Japanese style background music at least until someone loses the katana.

The game currently supports up to 16 players, and has a voting and a banning system. The game menu has filtering options and a very strong search feature, which can get a server list for any existing player name.

Being a free software project, Teeworlds has many custom maps and mods, a support forum on its website [www.teeworlds.com/forum], an IRC channel [#teeworlds on irc.quakenet.org] and is believed to have clans who battle out each other from time to time.

Oh, and as for my final words: this is an innovative, surprising and an addictive game that's easy to learn and real fun to play. 

By: Srikrishna Das

The author is a system administrator, toon artist, gaming addict and an engineering student. He is usually seen load balancing amidst his startup, mule in and college life. Canon fires or Gold coins can be sent to krish@deekix.com.

CONTENTS

ADMIN

- 56 Data Centre Redundancy
- 64 Wireless Intrusion Detection
- 68 The Art of Guard

GEEKS

- 58 Building a Server From Scratch—Part 3: Virtualising the Server
- 62 What Does Your Signature Mean?
- 88 Programming in Python for Friends and Relations—Part 12: Do You Speak my Language?

COLUMNS

- 51 FreedomYug: Orwell or Open-Source?
- 67 BraveGNUIndia: FOSS Campaigners Upbeat About Politicians' Stand on Software
- 87 The Joy of Programming: Understanding 'Side-effects' in C/C++ Programs
- 92 CodeSport
- 102 A Voyage to the Kernel—Part 11: Segment 2.4, Day 10.



REGULAR FEATURES

- 06 Editorial
- 08 You Said It...
- 10 Technology News
- 14 Industry News
- 16 Q&A Section
- 71 Linux Jobs
- 100 Tips & Tricks
- 108 FOSS Yellow Pages

Editor

RAHUL CHOPRA

Editorial, Subscriptions & Advertising

DELHI (HQ)
D-07/1, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: 011 26010602, 26010603
Fax: 26017563
Email: info@bityedia.com

BANGALORE
No. 19, 17th Main,
1st Cross, HAL, 8 Stage,
Indiranagar, Bangalore 560008
Ph: 0832 25260023, Fax: 25260094
Email: ctyb@bityedia.com

CHENNAI
M. Nockeeran
260 House, 31A, Cathedral Garden Road
Near Palmgrove Hotel, Chennai 600024
Ph: 044-28275191, Mobile: 09882500404
Email: ctych@bityedia.com

Customer Care

E-MAIL: support@bityedia.com

Back Issues

Kits 'n' Spares
D-07/5, Okhla Industrial Area,
Phase I, New Delhi 110020
Phone: 011 32975879, 26371665-2
Email: kits@bityedia.com
Website: www.kitsmagazines.com

Advertising

KOLKATA
D.C. Mitra
Ph: 033 22294788
Mobile: 22000094
Email: ctycal@bityedia.com
Mobile: 09433422932

MUMBAI
Ravi D'Souza
Ph: 022 24960047, 24938520, Fax: 24964278
Email: ctymun@bityedia.com

PUNE
Zahir Shabbir
Mobile: 09372407753
Email: ctypun@bityedia.com

HYDERABAD
P.S. Murthydharan
Ph: 09049962660
Email: ctyhyd@bityedia.com

Exclusive News-stand Distributor (India)

INDIA BOOK HOUSE PVT LTD
Arch No. 30, Indira Mahalaxmi Bridge, Mahalaxmi,
Mumbai - 400034 Tel: 24942530, 24925651,
24927353 Fax: 24950392
Email: info@ibhworld.com

Printed, published and owned by Rahul Chopra, Printed at Bityedia Office, C-118, DDA Flats, Okhla Industrial Area, Phase I, New Delhi 110020, on 27th of the previous month, and published from D-07/1, Okhla Industrial Area, Phase I, New Delhi 110020. Copyright © 2009. All articles in this issue, except for references, editorials, queries, or unless otherwise explicitly mentioned, will be released under Creative Commons Attribution-Share Alike 3.0 licensed license a month after the date of publication. Refer to <http://creativecommons.org/licenses/by-sa/3.0/> for a copy of the license. Although every effort is made to ensure accuracy, no responsibility whatsoever is taken for any loss due to printing errors. Articles that cannot be used are referred to the authors if accompanied by a acknowledgment and sufficient time and effort is taken for any loss or delay in returning the material. Elsewise, if any, will be settled in a New Delhi court only.



LFY DVD: Debian 5.0

OpenSolaris, a secure, stable and scalable operating system for users, developers and deployers, provides all the tools users expect from a modern computing environment—both installed by default and available online.



LFY CD: OpenSolaris 2008.11

OpenSolaris, a secure, stable and scalable operating system for users, developers and deployers, provides all the tools users expect from a modern computing environment—both installed by default and available online.

Notes: All articles in this issue, except for references, editorials, queries, or unless otherwise explicitly mentioned, will be released under Creative Commons Attribution-Share Alike 3.0 licensed license a month after the date of publication. Refer to <http://creativecommons.org/licenses/by-sa/3.0/> for a copy of the license.

ICFOSS

Wide Open Windows

The International Centre for Free and Open Source Software, which is to be set up in Thiruvananthapuram, is to be a leading research organisation.

The open source movement in Kerala is all set for a giant leap. The state government has taken a bold initiative to promote research in free and open source software (FOSS). Officials claim that nowhere else across the globe has such a programme been initiated under the auspices of a state government. ICFOSS, or International Centre for Free and Open Source Software, is to be set up at Kerala's capital city Thiruvananthapuram, and is envisioned to be a leading research organisation in the free and open source model of knowledge.

"We have a clear position on the politics of knowledge control and that goes well with the position of the free software movement. In order to further the objectives laid out in our policy, we decided to create an institution that can provide leadership and vision," VS Achuthanandan, chief minister of Kerala, told LFY.

Officials say this initiative, which was approved by the state cabinet late last year, will spread the idea of 'openness' in all aspects of society. Its focus is on directing the open education initiatives based on sharing knowledge and to explore the open source concept in hardware design. ICFOSS will identify and quantify the political, cultural and economic benefits for India as a result of FOSS. Arun M, the special officer for the centre, says: "ICFOSS is going to create an action plan to make India a global leader in the open source community. We will also leverage the open source development model based on community collaboration and shared ownership of intellectual resources to bridge the digital divide. We will initiate R&D in FOSS."

Arun says the centre will create a road map for open source software development for India's software export industry and will also provide infrastructure support for FOSS-based initiatives in different parts of the state. In sectors such as e-governance, education, local information publishing and in stimulating local employment, ICFOSS will pilot FOSS-based interventions.

Consultancy and certification

According to government officials, the centre is planning to work with local, national and international institutions and networks that operate in the FOSS domain. Chief Minister Achuthanandan, who also holds the IT portfolio, says ICFOSS will ensure Kerala's leadership in the FOSS sector. Though it will start out being a small group, it will have a global vision. It will help the state government by providing intellectual inputs in

creating initiatives, and by providing leadership to various free software and free knowledge initiatives in the state. Creative work will also form a part of the ICFOSS.


Of course, developing open source software is the core mission of the centre. But it will also address issues like patents, copyright, digital content, scientific publishing and so on. ICFOSS will help the government, public enterprises, educational institutions, small businesses, individuals and some key sectors in using IT. Ajay Kumar, secretary of IT, government of Kerala, says: "The centre will create infrastructure to support FOSS-based initiatives in different parts of the state to enable people to access IT support services, and will stimulate local employment through providing support services to FOSS-based Virtual Micro-Enterprises (VMEs), so as to enable them to start, sustain and expand operations."

The centre will also work as a consultancy to all government departments in Kerala in preparing Request-for-Proposals (RFPs), IT-related purchases, software acquisitions, certification of educational and training programs in open source, localisation of Indian languages, and to develop speech interfaces on FOSS. Open initiatives like open education, open hardware and open science initiatives will be taken care of by the centre.



Networking and research

ICFOSS will sponsor academic consultants and scholars to do research on the latest technologies in FOSS and network with leading R&D organisations in the same domain. Kumar says: "The centre will be in touch with the free software community all over the world. As a part of this mission, an International Conference on Free Software and Free Society was conducted in December 2008 in Thiruvananthapuram. Several government and non-government agencies worldwide and key players in the free software movement attended the meet. ICFOSS will build upon these relationships to make the community stronger."

The authorities are in touch with government and academic institutions from India and abroad. A delegation representing the government of Venezuela has shown interest in collaborating with the centre. On the wings of such lofty hopes, ICFOSS will start functioning by March 2009 -- but only if it doesn't get stuck in the 'politricks' the host state is notorious for. 

By: **KTP Radhika**

The author is a Linux admirer and correspondent at LFY's Bangalore Bureau.



Niyam Bhushan

Orwell or Open-Source?

Firmware is the next battleground.

Have you ever clicked a photograph from a mobile phone or a digital camera? Or listened to MP3 music in a car or home stereo? The next time you do, pause to look at the device at your fingertips, and tell yourself, "It's all about software."

Running inside the camera is the most sophisticated suite of software, working hard to conjure a perfect picture from the initial burst of pixels captured by the image-sensor. Almost all cameras provide embedded tools to touch-up and enhance your images. Think of it as a Photoshop or a GIMP inside the camera. Move over to your MP3 player in the car or the home-stereo. At the most basic level, there's MP3 decoder software at work inside the box. Not content with just that, most manufacturers have started to ship special software that analyse the MP3 stream and try to synthetically restore the original splendour of sound, lost while compressing the file into MP3. Similarly, a digital photo frame comes embedded with a JPEG decoder, and several other pieces of embedded software, based on the capabilities of the frame. In almost all cases, the embedded software is proprietary, and literally locked away in its sleek shiny box.

Consuming digital

Let's zoom out to look at the larger picture here: Consumer electronics is turning digital. Today, any digital consumer device ships with embedded software. The market for such consumer electronics (and if you add industrial devices) is many orders of magnitude bigger than the PC industry. In effect, each year, tens of billions of devices ship with embedded software that's all proprietary. Aggravatingly, ordinary appliances and machines are also beating with digital hearts. Take the cars I've owned over the last ten years. They've all had digital controllers embedded inside. Ditto for newer televisions and some washing machines. Even ordinary household switches are being transformed into digital home lighting solutions. Our lifestyle is obsessed with devices that promise us more creature-comforts, and even more creature-conveniences. Any device that brings us entertainment to dispel civilized boredom is always eagerly embraced with a fetish quite similar to a childish craze for newer toys.

Alas, the problems with proprietary software embedded so ubiquitously are age-old. Indeed, they may take on quite a sinister turn. Some manufacturers invoke the draconian

Digital Millennium Copyright Act (DMCA) if you try to hack your legally-purchased device. How would you know your MP3-player in your mobile phone does not report every unauthorised track that plays on it? Can you be sure the appliances you've bought with your hard-earned money, obey you, and are not remotely controlled and monitored by unseen masters? What about ATMs and voting machines?

'Begware'

Just wait. These fundamental problems are bound to explode rather dramatically in our world. When they do, people may finally realise the value and vision the free software community has been trying to highlight since the past 20

years. GPL version 3 partly tries to address the problem, thanks to its 'Tivoization' clause. Search the Web to familiarise yourself with this term if you don't really understand it. However, Tivoization does not solve the problem, as it still applies to free software that may get embedded.

Disappointingly, a large number of FOSS projects that may be used as embedded software have declined to use the GPLv3 so far. Leading the pack here is Linus Torvalds himself, with his Linux kernel. Even if the entire FOSS community were to adopt GPLv3, it won't solve the problem, since it's not necessary that all manufacturers embed FOSS in their devices. What's used inside all iPods, and the newer Amazon Kindle book-readers?

What's needed is a new approach that strongly advocates firmware-openness; and maybe even includes trade and government authorities to push for complete transparency. Until that happens, consumers may just find themselves reduced to 'Begware'. Consider Apple iPhone users, paying a premium amount for a fancy phone, and then begging Apple to include features they want, or software they wish to buy, from Apple's Web-store. Just think: if you can't even be the master of any device you own, we might just be living in a new digital world of insidious slavery. 

About the author:

Inspired by the vision of Osho. Copyright February 2009: **Niyam Bhushan**, freedomyugs at gmail dotcom. First published in LinuxForYou magazine. Verbatim copying, publishing and distribution of this article is encouraged in any language and medium, so long as this copyright notice is preserved. In Hindi, 'mukt' means 'free-of-cost', and 'mukt' means 'with freedom'.

"The problems with proprietary software embedded so ubiquitously are age-old. Indeed, they may take on quite a sinister turn."



Why Not Carry A Pen Drive Instead?

Carry it wherever you go. Boot from any computer. Have your data along with personalised system settings available anywhere, with negligible additional weight in your pocket... That's what Linux on your pen drive offers.

*I*n this tutorial, we will install five different Linux distributions on a pen drive. My shortlist is Slax, Damn Small Linux (DSL), GoblinX, Puppy Linux and Pen Drive Linux. Of course, you can choose to have something else.

Let's start off with the requirements:

- A Fat16-formatted pen drive. Except for Puppy Linux and DSL, which would happily install on a 512 MB drive, I would recommend a pen drive of 1GB or more. Mine is a 4 GB (U3 removed) pen drive.
- ISOs of Puppy Linux and DSL, the tar package for Slax, GoblinX and Pen Drive Linux zip packages for USB. Download the latest versions for up to date software.
- A motherboard BIOS that is capable

of booting from USB drives. Most motherboards from 2004 onwards have this feature. This is not necessary if you decide to use PuppyLinux.

- Any version of Linux installed. I use openSUSE 11.

Points to remember before getting started:

- Before using your pen drive, back up whatever data is in it, as it may get deleted.
- If you decide to install on a U3 drive, you must remove the U3 software before you decide to install. Check out U3 removable tools from www.u3.com/uninstall.

Getting your pen drive ready

Your pen drive needs to be Fat16 formatted. Use GParted, *cfdisk*, or any other partition tool for the purpose.

First, you need to identify your pen drive with *fdisk*. Plug in your pen drive and run the following command as the root user:

```
fdisk -l
```

This should show your pendrive as sdX, where X is a, b, or c... Make a note of it.

Follow this step (for *cfdisk*) only if you need to format your drive as Fat16—remember that all your data inside the pen drive will be lost:

```
cfdisk /dev/sdX
```

Choose 'Type' as 06, 'Bootable' and 'Write', followed by *yes* and 'Quit' to write to the partition table of the pen drive. Note that the sequence of what you need to do (as mentioned here) is important.

Slax on Pen Drive

Download the Slax 6.0.9 tar package for the USB from www.slax.org/get_slax.php. Now, carry out the following steps:

1. Double click on the Slax-6.0.9.tar package downloaded. This will open the Archive Manager and show the contents of the package, namely */boot* and */slax*. Drag and drop them onto the USB partition.
2. Open the terminal and navigate to the */boot* directory copied on the USB. You will see a shell script called *bootinst.sh*. Run the script as the root user:

```
sh bootinst.sh
```

3. That's it! You now have Slax on your pen drive!

Remarks: Very easy to install. Good selection of packages. Easy to add modules to the distro—just download the required modules from the Slax website and use the Slax Module Manager. Good hardware detection. Persistence settings.

Damn Small Linux

Download the ISO from www.damnsmalllinux.org/download.html and burn it on a CD. Boot the live CD, and plug in your pen drive. Now follow the steps given below:

1. Navigate to *Start→Apps→Tools→USB-HDD Pen Drive Install*
2. In the terminal that opens, you will be asked to answer a couple of questions:
 - Verify that your USB device is sdX. Answer 'y' if this is the correct device as shown by *fdisk* earlier.
 - The boot options will be presented next. Press *Enter* if they are acceptable.
 - Choice of keyboard: Answer 'us' for the US keyboard and press *Enter*.




Figure 1: Installing Slax on the USB

- Following this you will be asked whether you are ready to proceed with the install. Press *Enter*.

DSL will create two partitions—sdX1 (a F16 partition) and sdX2 (an ext3 partition)—and install the GRUB bootloader.

Remarks: Easy to install. Very good selection of packages. Additional packages can be downloaded by first upgrading to GNU Utilities (*Start→Apps→Tools→Upgrade to GNU Utilty*) and then enabling Apt. Required applications can then be downloaded from *MyDSL→My DSL Browser*. Excellent hardware detection (borrowed from Knoppix). Persistence setting can be enabled from *Start→Tools→Control Panel→Backup/Restore*, and entering sdX1 or sdX2 as the partition to save system settings in the window that opens. The only drawback is that some of the packages from repositories don't work and cause the system to freeze.

GoblinX

This is one of my favourites. Download the GoblinX-2.7.zip package for the USB from www.lax.ic.unicamp.br/pub/goblinx/GoblinX-2.7/GoblinX-2.7.zip.

To install on your pen drive, follow the same steps carried out for Slax:

1. Double click on the GoblinX-2.7.zip package. This will open the Archive Manager and show the contents of the package, namely */boot* and */goblinx*. Drag and drop them onto the USB partition.
2. Open the terminal and change to the */boot* directory copied on the USB. You will see a shell script called *bootinst.sh*. Run the script as the root:

```
sh bootinst.sh
```


Remarks: Very easy to install. Excellent selection of packages. Very easy to install additional packages through Gslapt. Good hardware detection. Persistence settings. The Mini Edition has a choice of various desktops.

PuppyLinux

This is a 90 MB distro that packs quite a punch with numerous utility packages. Download the ISO from www.puppylinux.org/downloads/official-releases/puppy-linux-412 and burn it to a CD.

Boot the live CD, plug in your pen drive and follow these steps:

1. Click **Start→Setup→Puppy universal installer** and choose the USB flash drive install from the wizard that follows. Your pen drive will be detected and you will be presented with a choice of two methods to install Puppy to the USB:
 - A Combo Format method, which creates a 128 MB F16 partition for Puppy system files as well as Pup save configuration files, and a large ext3 partition for data. I found that this method only works if you don't save persistence settings, as the 128 MB size is just too small. I wasn't able to save persistence settings to the ext3 partition.
 - The Super Floppy mode simply takes you to Gparted.
2. To install Puppy, click on the **Install Puppy to sdX** button located on top of the Puppy Universal Installer wizard. See Figure 2. You will then be asked if your PuppyLinux CD is in your CD-ROM drive. Press **Enter** if it is, and then choose to have syslinux installed from the next installer wizard page. An orange box will then show up on the desktop. Now you need to just press the **Enter** key for the 2 - 3 questions asked.

 **Note:** Do not use the partition tools found on PuppyLinux. In my case, Puppy Universal Installer detected the file system wrongly when created with these tools. If your pen drive is already F16 formatted, refer to Point (2) above. If not, use partitioning tools like Gparted or *cfdisk* from the Linux distribution that's installed on your hard disk, and then follow instructions in Point (2) above.

Remarks: Not easy to install. The partition tools made a mess of the filesystem. Does not install to an ext3 filesystem, in spite of the Universal installer claiming to be capable enough. Excellent selection of packages. Very easy to install additional packages. Systems administration is a cake walk. Excellent hardware detection. Persistence settings. Fabulous wallpaper and a light desktop (JWM). Download the boot floppy package from the Puppy Linux site for motherboards that do not support direct booting for the USB.

Pendrive Linux

This distribution is based on the discontinued MCNLive, a subset of Mandriva. Download the USB package from downloads.sourceforge.net/pendrivelinux/Pendrivelinux08.zip


To install, carry out the following steps:

1. Double click on the Pendrivelinux08.zip package.




Figure 2: Puppy's Universal installer wizard

- This will open the Archive Manager and show the contents of the package. Drag and drop them onto the USB partition.
2. Open the terminal and change to the */boot* directory (see note below) copied on the USB. Run the *bootinst.sh* script as the root user. Login password for root is 'root' while for the guest, it is 'guest'.

 **Note:** Pendrive Linux can only be made bootable through Windows (ugggh!). To avoid using Windows to make the drive bootable, I copied the */boot* folder of Slax (or Goblinx) to the pen drive and then executed the *bootinst.sh* script. After that, I deleted the */boot* folder.

Remarks: Very easy to install. Excellent selection of packages and hundreds more available through Mandriva repositories. Very easy to install additional packages. Good hardware detection borrowed from Mandriva. Persistence settings. (Default loop size is 256 MB. Larger loop size can be created by navigating to *Start→PendriveLinux→Create persistent loop*, which only works when you choose the bootloader option 'Boot Live'. Once created, you need to use the 'Boot with Persistence Changes' bootloader option to keep your changes.) Systems administration made simple through Mandriva Control Centre. 3D Desktop (Beryl and Compiz) choices make this distribution a frontrunner for the USB.

Bottom line

Having a hard time choosing the best distribution for your pen drive? Well, I, too, am facing the same dilemma. 

By: Nelson Lobo

The author is a lecturer at St Xavier's College, Mapusa, Goa. He can be reached at lobonano@gmail.com

Your Daily Dose of TECHNOLOGY NEWS



IF ITS TECHNOLOGY IT'S HERE

**Free
Subscription**

Simply visit www.efytimes.com
and **subscribe** to the newsletter

DAILY TECH UPDATES >>>

From the EFY Group on subjects ranging from nanotechnology to open source software can land in your mailbox for **FREE!**

Hardware, Software, Consumer Electronics, Robotics, Automotives, Telecom, Space—any news that's related to technology will be featured here.

www.efytimes.com



EFY Enterprises Pvt Ltd

D-87/1, Okhla Industrial Area, Phase I, New Delhi 110 020; Ph: +91-11-26810601-03
Fax: +91-11-26817563, 26812312; E-mail: info@efyindia.com; Website: www.efyindia.com

Data Centre Redundancy

If you have been following this series, you must have set up a highly available and load-balanced cluster of Web servers. As discussed in the previous article, we needed data centre redundancy in order to sustain site failures. In this article, even if we don't cover the perfect architecture, we hope to give you some insights into how this objective is attained.

The primary objective of this article will be to show how to serve the website <http://unixclinic.net> from two different data centres, where the load is distributed fairly between both the centres. While reading the description, it is advisable to go through the layout in Figure 1, which basically gives the architecture we'll discuss in the following sections.

Layer 0: DNS round-robin load balancing

We have discussed DNS round-robin load balancing in the first part of the series (February 2009). This is our first level of load balancing. In this, typically, we need to create multiple A records for *unixclinic.net*, each having a different IP address. Each client accessing the <http://unixclinic.net> website will get a different IP address for it by the DNS in a round-robin fashion. This is not the perfect load balancing, because there is caching involved in the DNS hierarchy and the browser also. Due to this

reason, it is important to have a decent TTL value for the DNS record.

Typically, there should be one A record in the DNS, per data centre.

Once the client gets the IP address, it will try to access port 80 on that IP address. The IP address will actually be owned by the external firewall in the data centre. So, at this stage, the user can either reach data-centre-1 or data-centre-2. Let us assume that this connection has accessed the external firewall at data-centre-1.

Layer 1: External firewall (external zone)

The traffic for *unixclinic.net* will be received by the external firewalls in each data centre. These will be high availability firewalls. This firewall will also double as a load balancer to balance traffic between data centres. So we need to choose a product carefully, one that can provide the functionality of a firewall as well as a load balancer. Given a choice, I would rather go for Linux Netfilter/iptables with

customised rulesets for firewalls, LVS (Linux Virtual Server) for load balancing and Heartbeat for HA needs.

Whatever solution you choose, the firewall will be set to allow incoming port 80/tcp connections (and port 443/tcp if you are serving some secure content through SSL), and redirect them to the reverse proxy servers on the DMZ, which are running nginx.

The external firewall on data-centre-1 will load balance the traffic between the reverse-proxy server on each data centre. The algorithm for load-balancing should be wisely chosen. LVS supports round-robin, least connections, weighted round-robin and weighted least connections. Avoid round-robin here for two reasons: one, so that there are more chances of connections being distributed evenly and two, we are already using round-robin load-balancing at both DNS and reverse-proxy levels.

A connection that reached data-centre-1 may either go to reverse-proxy servers located at data-centre-1 or at data-centre-2, depending on the decision taken by the load-balancing algorithm.

Layer 2: nginx reverse proxy servers (DMZ)

The packet redirected by the external firewalls may reach either of the reverse-proxy set-ups at each data-centre. Both the reverse-proxy set-ups are configured identically and contain all the six nodes (Web-servers) as their members. So, irrespective of which reverse-proxy the connection reaches, it may actually be processed by any of the six nodes.

In the first article in this series, we explored the set-up of the nginx reverse-proxy, so I will not be going into too much detail here.

Layer 3: Web servers (internal zone)

The Web servers may run either nginx or Apache, depending on the requirement. Again, we have already covered the configuration of nginx as a Web server with PHP support in the first article itself, so details are not required.

One of the important points here is the website being served. Since quite a lot of the code will be the same, should we copy the same code across all Web servers while deploying the code or should we make use of a network-based common storage pool like NFS or SAN? The answer to this is based on your requirement and whether or not you would like to add a few more blocks in the architecture. For a website, static or dynamic, I would prefer to write a deployment script that can deploy the code to all the Web servers and not use SAN or NFS.

The second important point is the database. If you are serving a dynamic website that uses databases, then you should ideally configure the Web servers located in each data-centre to connect to the database situated locally. This will save the slow SQL queries to be executed over a




Figure 1: The architecture

WAN link between data centres. The application utilising the database should typically use a connection pool including the database at both the data centres, with a higher priority for the database situated locally.

Layer 4: Databases (internal zone)


The database layer in our architecture is purely optional, but if you do have one, then there are certain points you need to keep in mind. The database queries should only be using the database located in the same data centre where possible, and only in case of its failure should they go to the other data-centre.

Typically, for this kind of set-up, there will be one database in each data centre with a two-way replication set-up between them.

Moving further

In this article, we have seen how we can achieve data centre redundancy to increase the availability of your website.

For the geek in you, I have given enough pointers to get you started. I sincerely hope that you have enjoyed these three articles and will take advantage of them to put this architecture into production, to demonstrate the benefits of using FLOSS to your organisation.

This is definitely not the last of the articles in this series, as we are yet to cover the firewall configuration and data centre load balancing using LVS. We also need to look at setting up a database cluster with two-way replication. This will be covered in forthcoming issues. 

By: Ajitabh Pandey

The author has more than 13 years of diversified IT industry experience in training, support and consulting. You can know more about him at ajitabhpandey.info and e-mail him at [ajitabhpandey \[at\] ajitabhpandey \[dot\] info](mailto:ajitabhpandey[at]ajitabhpandey[dot]info)

BUILDING A SERVER FROM SCRATCH

Virtualising the Server

Set up VirtualBox on the main server and deploy each service on a separate virtual machine.

I know, I know... you were expecting how to configure a Web server this time. However, the set-up that I described has more than a few glitches. First of all, putting in all the software in one environment can lead to data getting polluted. Second, since I was going to do virtualisation anyway, why just keep it as a role? Why not put it to use? So, anyway...

Pre-summary

- Role: Virtualisation
- Priority: Required
- Dependencies: Basic server
- Features: A VirtualBox server environment with one virtual machine (VM) dedicated to each role. Each VM will be running CentOS 5, a free-of-cost

version of Red Hat Enterprise Linux (RHEL) 5 built with the same sources.

- Extensible: Yes
- Ease of set-up: Very easy but needs a lot of patience.
- Software Set: Sun xVM VirtualBox, CentOS 5.2
- Set-up time: 3 hours

A little explaining

Let me explain what we are going to do.

Virtualisation is all about packing more into one. Basically, the kind of virtualisation we are going to do here allows us to install more than one operating system on one computer and then run them simultaneously. For that, we need a hypervisor—a software that manages this virtualisation. Our

hypervisor of choice here is the Sun xVM VirtualBox.

Beware! Most distros pack in VirtualBox OSE. The OSE here stands for Open Source Edition. VirtualBox has only been partially open sourced. Some of the advanced features have not been open sourced yet and are not available in this OSE. One of these features is VRDP support, which we need for remote access to our machines. So, we have to use the free (as in free beer) but closed source version of VirtualBox.

Some of you may ask, if we have to use closed source, why not go for VMWare Server? It's free and it's made for remote access and administration. We are going to use VirtualBox because VirtualBox is faster and more stable on Linux. Additionally, part of VirtualBox is open source. More specifically, the core engine is open source; only the VRDP bit is closed source.

This brings us to remote access and administration. Unlike VMWare Server, VirtualBox is not specifically meant for servers. It's a general hypervisor. So it has nothing like VMWare's Web Interface. On the other hand, it has a very powerful command line interface called VBoxManage, which actually exposes more features than the VirtualBox GUI. So you can SSH into your server from remote terminals and configure virtual machines. And since SSH allows you to run X applications as well, you can use the GUI to configure VMs too.

What about viewing machines? Well, that too is simple. I've been ranting all along about VRDP. VRDP (VirtualBox Remote Desktop Protocol) is VirtualBox's implementation of the Microsoft Remote Desktop Protocol implemented after it had been opened up via MSDN. This feature allows you to use a RDP client like Microsoft Remote Desktop Connection (bundled with Windows) or a desktop on Linux to see a VM's Desktop.

Now about the guest operating systems...everyone acquainted with servers should know all about RHEL, right? Well, the fact of the matter is, it's extremely expensive. The CentOS Project (Community Enterprise Operating System) makes use of the free RHEL sources, replaces the RHEL branding with CentOS branding and thus produces completely free binaries, which are a 100 per cent binary compatible with the corresponding RHEL version, thus creating a free RHEL clone.

The CentOS project even publishes the security updates. So, users of CentOS are technically using RHEL, just re-branded to escape licensing issues. By the way, RHEL and CentOS are completely separate entities and are not affiliated to each other in any way.

So let's get started, shall we?

Preparations for the launch

We need to prepare for our mission. First, there are some prerequisites that you have to set up yourself. Make sure you have a Linux or Windows (preferable) client attached to your server, which can browse

the Internet. This is to make sure that the client can communicate with the server, since all Internet bound traffic will have to be NATed by the server. Ping the server if you need more reassurance. Traceroute it if you need still more. I asked for a Windows client to be able to demonstrate the interoperability of the solution, but if you use Linux, it's just as well. I'll still give you instructions, don't worry.

Then, since we were going in a different path earlier and have changed tracks now, you need to change the partitioning scheme. Create a single LVM volume out of all the unused space of your HDD and mount it under */VMspace*. That should do.

Now, let's do some work.

Installing and configuring VirtualBox CSE

Sadly, CSE stands for closed source edition. But, there's nothing we can do about it. First of all, we need to remove VirtualBox OSE if it's there. You can use Synaptic to do it.

Then, run the following command to clean up any stale dependencies:

```
# sudo apt-get autoclean
```

Now point your browser to www.virtualbox.org/wiki/Linux_Downloads and download the VirtualBox package for Ubuntu 8.10 (remember this is the host OS we've installed on your server). Once done, install it:

```
# sudo dpkg --get-selections virtualbox-2.1.2 | sed -e 's/hold/hold hold/'
```

After your installation, issue the following command to resolve all dependencies:

```
# sudo apt-get install -f
```

By the way, here's one tidbit: the closed source VirtualBox is free for personal use. This may scare you because enterprise use isn't exactly personal, is it? However, take a look at the VirtualBox Licensing FAQ Point 6, and all your fears will be put to rest. Basically, this is what it says:

"Personal use is when you install the product on one or more PCs yourself and you make use of it (or even your friend, sister and grandmother). It doesn't matter whether you just use it for fun or run your multi-million Euro business with it. Also, if you install it on your work PC at some large company, this is still personal use. However, if you are an administrator and want to deploy it to the 500 desktops in your company, this would no longer qualify as personal use. Well, you could ask each of your 500 employees to install VirtualBox but don't you think we deserve some money in this case? We'd even assist you with any issue you might have."

Editorial

Dear Readers,

I must start off by saying, "Thank You, Chennai!"

We are lost for words to express our gratitude for the amazing response and support we received for Open Source India 2009. This was, by far, the most successful edition of Open Source India (and LinuxAsia) that we have organised in the past six years!

Every year, once the event gets over, we review the highs and the lows. This year, the list for 'highs' was long, while that of 'lows', much shorter. What topped our list of 'highs' is getting the target audience, and in appreciable numbers.

Right from Day One, we had decided to focus on different user segments of Open Source.

First, we targeted those who take the strategic decisions and set budgets—the CXOs. From our last year's experience, we had realised that virtualisation was a topic that many from this segment were interested in. So we created a CXOSummit focusing on virtualisation, cloud computing and SaaS. Now getting the busy senior executives in this segment is very difficult, but at Chennai, we achieved a 'House Full'.

We then looked at the IT managers and systems administrators, who actually implement IT in their organisations. In fact, for small and medium firms, they often play the strategic role of CIOs too. We had a full day session for them, where they were to be presented with Open Source options for various tools and resources required for IT management. The response? We had to shift the session from the 120-seat hall to one that could at least accommodate 200. The final attendance was over 230.

So, we had the decision-makers and influencers covered. Next, we focused on those who would develop applications for them—the developers. We had two seminars for them—one focused on software development for mobiles and the other for everything else. We also opened these sessions for budding developers—the

engineering students—so that they were exposed to and hence motivated to adopt Open Source. The seating capacity for this event was 500. At times, the hall witnessed over 800 attendees!

Last, but not the least, were the members of academia who were on our list. If the next generation is to be open source aware, then it's important that their *gurus* are convinced about its merit too. We organised a special half-day workshop for them. We'd planned for 50, and finally welcomed 80 into a 70-seater hall.

So we managed to get quite a few people in, but were we able to deliver the right message? First hand feedback at the event was positive. But we decided to not go by our individual perceptions, but to poll our visitors. At the time of writing this piece, we had received feedback from 350 attendees through our website (www.osidays.com) and on an average, we had scored a 'Good'.

So we managed to get quite a few people in, but were we able to deliver the right message?

While there are many who need to be thanked for making OSI 2009 successful, someone who tops the list is Dr Rajagopalan, who heads C-DAC, Chennai. He had repeatedly told us that Chennai was the city that deserved an event like Open Source India. Now, we hope to keep OSI at a level, where it lives up to the standards set at Chennai.

Best Wishes!



Rahul Chopra
Editor, LFY
rahul@fyndia.com

Now, add the root user to the *vboxusers* group:

```
# sudo usermod -G vboxusers root
```

Since we are using VirtualBox 2.1, we don't even need to create a network bridge. All we do is plug in our VMs to *eth1* (that's our client interface), and yippee, we are networked. *eth1* has DHCP, so the virtual-servers are automatically configured. A word of caution: since we have a DNS system in our network, all our VMs should have good host names. I don't know what you named your domain, but suppose it's *officenetwork.local*, give your database server the name *db-server* and the mail server the name *mail*. So we can access them as *db-server.officenetwork.local* or *mail.officenetwork.local*.

Note that our architecture does not allow for accessing any of our VMs from outside (more precisely, *eth0*) our office. So a field officer trying to access his mail, will need to be able to VPN into our server. I'll set this up at a later stage.

Setting up our first VM: The Web server

First, some disappointing news... We will not set up our Web server now but just prepare it.

VirtualBox 2.1 can run 64-bit guests on a 32-bit host and vice-versa. So, since you have a Q6600 processor, you can seriously consider a 64-bit version of CentOS. And your processor supports hardware assisted virtualisation, so it's all the better.

First of all, head to *centos.org* and get yourself all the 5 CDs or the DVD (torrent download only) of the latest version of CentOS. As of now, CentOS 5.3 (corresponding to the RHEL 5.3 release) has not been released, so you'll have to settle for 5.2. However, CentOS 5.3 may come out any day now. You can upgrade as per your needs, easily.

I assume that you have downloaded the CentOS 5.2 x86_64 DVD and have saved it in */VMspace/ISOStore*.

So open up a terminal and type in the following:

```
# sudo mkdir -p /VMspace/VDisStore
# sudo VBoxManage createvm --name 'WebServer' --baseid 1
# sudo VBoxManage createvdi --filename \
  /VMspace/VDisStore/WebServer.vdi --size 5122 --format
# sudo VBoxManage modifyvm 'WebServer' --id 1 \
  /VMspace/VDisStore/WebServer.vdi
# sudo VBoxManage modifyvm 'WebServer' --cd '( \
  /VMspace/ISOStore/CentOS 5.2 x86_64 -bin-DVD.iso)
# sudo VBoxManage modifyvm 'WebServer' --memory 512M
# sudo VBoxManage modifyvm 'WebServer' --audio -- \
  --audio on --audio on --audioport 5120
# sudo VBoxManage modifyvm 'WebServer' \
  --ostype RedHat_64 & change if using 32bit
# sudo VBoxManage modifyvm 'WebServer' --passthrough \
  --virtex on --acpi on --hpet on
```

Now, to configure some sophisticated networking. This is possible from the *VBoxManage* command line as well, but to better understand the procedure, let's do this from the GUI.

Launch the GUI, select *Web Server* from the VM list and then click on the networking hyperlink on the right. In the pop-up box, enable Adapter 1. Then, select its type to be "Intel Pro/1000 T Server (82543GC)". Choose 'Host Interface' in the *Attached To* field. Generate a random MAC address (do not change this in the future) and finally attach the host interface to *eth1* (our internal network).

So NASA has just given us the Go-code for the launch! However, you soon abort the launch after noticing a fatal error—it's impossible to access the Web server from outside (*eth0*). Well, relax! What are reverse proxies for, anyway? By the way, give *nginx* a miss; we will set up a different configuration another day.

Fine-tune any settings as you see fit (if you need audio, you can set it to AC'97) and then drop down to the command line. Go to your Windows machine and start up the Microsoft Remote Desktop Connection. On the address field, use the server FQDN (server.officenetwork.local, for example) and append a ":3390" to it. Expand the options button, click on the experience tab and select "LAN (10 Mbps or higher)". Don't connect as yet.

Type in the following command at the server console:

```
# sudo VBoxManage startvm 'WebServer' --type vnc
```

You should get a pop-up window running *VBoxHeadless* with the VM. Immediately, connect the RDC on the Windows machine. Voila! You have the Remote Access part of VirtualBox running as expected. Don't boot up now, abort the GRUB countdown and issue the following:

```
# sudo VBoxManage stopvm 'WebServer'
```

Launch up the GUI, and continue installing CentOS 5.2. Why I told you not to set-up remotely is because X behaves radically if you don't have the VirtualBox Guest Additions installed.

At this point don't set up any server tools at all—just set up a basic X desktop environment, preferably GNOME (because it is lighter). Just have one partition (and a swap-space), no LVM (the KiSS strategy). Give it a descriptive hostname ('web' would do fine) and set it up to configure its network over DHCP. Oh, and disable SELinux and the firewall in the text-mode part that comes after Anaconda's finished.

Install all the VirtualBox Guest Additions now. RDP into the machine (no problem if the GUI is running as well; anyone sitting at the GUI will think that a ghost has possessed the VM) and see if it works well. It should.

A SSH server for remote administration

Since VirtualBox does not have a dedicated remote administration tool, we will have to SSH into the machine and use *VBoxManage*. So, we need to set up an SSH server in our server.

First install the OpenSSH server:

```
# sudo apt-get install openssh-server
```

Now open the */etc/ssh/sshd_config* file and make a few changes: Set the port to something non-standard, such as 99; set *RSAAuthentication* to *Yes*; set *Protocol* to 2; *PasswordAuthentication* to *NO* (we will set up certification-based authentication); and *PermitRootLogin* to *Yes*.

Now we need to set up Public Key authentication. SSH keys are machine-specific, not user specific. So you need to set up one key on each machine that you want to access the server from. From a terminal on a Linux client, type:

```
# sudo mkdir -p ~/.ssh && chmod 700 ~/.ssh
# sudo ssh-keygen -t rsa -b 4096 -f ~/.ssh/id_rsa
```

Supply a really big passphrase. It can include anything. Just remember this 30 character (optimal) long sentence.

Then secure the files:

```
# sudo chmod go-rw ~/.ssh
# sudo chmod 700 ~/.ssh
# sudo chmod go-rwx ~/.ssh/*
```

You subsequently need to change two lines in the client */etc/ssh/ssh_config* file:

- Set *IdentityFile* to *~/.ssh/id_rsa*
- Set *Protocol* to 2

Repeat this for every Linux client.

In your clients, there will be a *~/.ssh/id_rsa.pub* file. This is the public key file. Copy these files, one by one, to your server and append them to the *~/.ssh/authorized_keys* file (use the *cat* command with the *>>* directive).

Change the server configuration a bit; in the *sshd_config* file set:

- *Protocol* to 2
- *PubKeyAuthentication* to YES
- *AuthorizedKeysFile* to *.ssh/authorized_keys*
- *PasswordAuthentication* to no
- *ChallengeResponseAuthentication* to no
- *UsePAM* to no

Now do a */etc/init.d/ssh* restart and you should have a fully working SSH environment. Try to log in, with the *"PreferredAuthentications=publickey"* flag to *ssh*. You should have it.

For more, consult stef.org/howto/openssh/publickey-auth.

Appendix A: RDP client on Linux

The RDP client on Linux is an X application known as *rdesktop*. To install it, use:

```
# sudo apt-get install rdesktop
```

...or, on your CentOS VMs:

```
# yum install rdesktop
```

You can connect with one simple command:

```
# rdesktop -u user -p password -i local 2390 # adjust the username and port
```

For more information, see the man page.


Appendix B: SSHing from Windows

You use a free program called PuTTY for Telnet and SSH on Windows. PuTTY can be downloaded from tartarus.org/~simon/putty-snapshots/x86/putty-installer.exe.

Install this snapshot. There are a few guides to using PuTTY on the Internet. Use the following two websites to work out how to configure Public Key authentication on PuTTY.

- www.ualberta.ca/CNS/RESEARCH/LinuxClusters/pka-putty.html
- searchenterpriselinux.techtarget.com/tip/0,289483,sid39_gc11272986,00.html
Use a 4096 bit SSH2-RSA key pair.

Some notes

1. You must specify a separate VRDP port for every VM. You cannot remote access VMs by name. They must be accessed by port. So maintain a list of names to port number assignments, somewhere.
2. You can use SSH to launch an X application remotely also (you need to use Linux here). Use this link: nixcraft.com/getting-started-tutorials/170-run-remote-x-applications-over-network-using-ssh.html.
3. This one is for creating purely random passphrases and passwords by rolling a dice a number of times: world.std.com/~reinhold/diceware.html. This is highly recommended for security. 

By: Boudhayan Gupta

Boudhayan is a 14-year old student who suffers from an acute psychological disorder called Distromania. He owes his life to Larry Page and Sergei Brin. Apart from that, he enjoys both reading and writing, and when he is not playing with his Python :-), during most of his spare time he can be found listening to Fort Minor, or cooking.

What Does Your

Signature Mean?



Ever wondered what those cryptic characters that some geeks use as their signatures, mean? I like to call them command-line gizmos.

I hang out a lot on *usenet*, mostly on *comp.unix.shell*, a group that has educated thousands of shell users for many years. Many of the posters here have signatures that are small scripts. Deciphering these scripts is very entertaining and pushes you to the man pages.

So I also started sporting these signatures. One good thing is people stop by and ask me about that cryptic text I sport at the end. When I explain, they get quite pepped up about UNIX/Linux.

My current signature is:

```
unset BC_LINE_LENGTH=10000
awk BEGIN { for (i = 1; i <= 999; ++i) printf "%d\n", i }
-i 10000
perl -le 'print 0 + ($ARGV[0])'
```

So what does this script mean? Let me explain. An interviewer once asked me the following question.


Q. $x = 1*2*3*...999$

How many trailing zeroes will you get in the resulting number, x ?

And being groomed in the UNIX way of life, I wrote the script I've mentioned above, and told him we can find out the number of zeroes *this way*...

He liked it (though this was not an expected answer) and moved on.

When I came back home, I tried the same script out for a much larger range, up to 9999, liked it and made it my siggie. I also realised I was using some common UNIX command line idioms.

 **Note:** For newbies, all the above commands in the 'pipeline' are 'filters'—programs that read from standard input, do something useful and write it to standard output. You can run all these filters separately—as *simple commands*—so you can experiment with sample keyboard input and screen output. Filters don't care how they are used. It is the shell that arranges things; so filters, *unknowingly*, become part of a pipeline and without knowing it, talk to each other.

Going back to my signature – let's look at the second line:

```
awk BEGIN { for (i = 1; i <= 999; ++i) printf "%d\n", i }
-i 10000
```

In *awk*, *BEGIN* block executes as a preprocessing step... before any input is



WIRELESS WIRELESS Intrusion Detection

Wireless networks are more prone to unauthorised intrusions compared to wired networks. Are you prepared for these threats?

Wireless technology cannot be ignored any longer. Given its advantages and the conveniences that come with it, it will probably replace all current network infrastructures pretty soon. As with all new things, this technology, too, brings with it a number of potential loopholes. That's why we need to be more careful about guarding these networks from security threats.

Wireless networks are more prone to attacks when compared to wired networks. You do not have to physically connect your computer to the network to get access to a wireless network, which means that there are more chances of getting your network compromised. This makes securing a wireless network a tough job. An administrator of the network may use strong encryption and prevent the possibility of an attack to some extent, yet the use of encryption such as in WEP is not good enough and can be cracked very easily.

WPA and other new standards like WPA2 are stronger, but are still more vulnerable to attacks than traditional wired networking structures. What most common users or new users of wireless networks do not understand is, even the strongest encryption is useless if your passphrase is simple. Using passwords like "1234" or "abcd" with strong protocols is as good as no encryption at all because it can be cracked very easily. Even with the use of stronger passwords, the advances in cracking means that if a cracker tries hard

enough, he will eventually succeed in breaking through security.

With a majority of networks (all the discussion henceforth relates to wireless networking unless mentioned otherwise) that apply simple passwords, the possibility of potential crackers entering a network is increasing day by day. This makes the detection of network intrusion a very difficult task indeed.

There are some types of attacks that are initiated from rogue access points in the vicinity of a wireless network. Such attacks have nothing to do with the encryption protocol being used. As already stated, the inherent vulnerabilities of ad-hoc networks with simple passwords, makes you a very good potential target for crackers!

An ad-hoc network is essentially one that does not have any fixed structure. The literal meaning of ad-hoc is 'for this purpose only'. What happens practically is, some nodes come into the network only to do some specific task. Devices are part of the network only for the duration of a communications session till they complete that specific task, after which they leave the network—there is no rigid structure.

Need for intrusion detection systems

Intrusion prevention measures, such as encryption and authentication, can be used in ad-hoc networks to reduce intrusions, but they do not ensure total elimination. For example,

encryption and authentication cannot defend your network against clients that have already gained access to the network (compromised nodes or rogue access points), which often carry the private keys and other sensitive information.

Any structure or network is only as strong as its weakest link. Even after application of all security measures, we cannot safely say that the threats have been eliminated. Intrusion detection is an add-on security to the networks that help in forming the second line of defence against the potential attacks, and is an important addition to the networks that want to stay secure.

What's an Intrusion Detection System?

In simple words, an IDS or an Intrusion Detection System, is software that is used to detect unauthorised access attempts to a computer, mainly through network access. IDS does not, or rather, cannot prevent an attack, as many misguidedly believe. IDS only detects an attack being made in real-time. This means, when an IDS detects an intruder, some part of the network has already been cracked into and the attacker is already a part of the network.

Similarly, it is also not a firewall. It will not block traffic into your network. Though one can program the IDS to take some specific action if an intrusion is detected, such actions will be taken by some other application working together with the IDS.

A network IDS will continually monitor packets on a network and attempt to discover whether a system breach is in progress. Loss of information is not the only problem in case of wireless intrusion (or any network intrusion, for that matter). Some attackers may try to freeze your system, launch fork bombs or try to cause DoS (Denial of Service) attacks on the system.

Some IDS systems can be tweaked to do some specific tasks upon finding an intrusion, such as finding the IP address and blocking it by invoking the firewall or finding the MAC address of the rogue machine and tracking its location. What must be noted here is that both the IP and MAC addresses can be spoofed easily.

So given that this is only a monitoring system, one needs to constantly monitor the network to see if an intrusion is happening. It is like having any other security system. You may fit a camera or a tracking device at your home to monitor unauthorised activity, but unless you actually see the activity in real time, you cannot prevent it. Watching recorded tapes is only the evidence that something bad has happened. You cannot go back in time and undo the damage. Your valuables are gone.

How wireless IDS works

As already discussed, a wireless network is more prone to attacks because it eliminates the need to connect a computer physically to an access point. This is the main reason why wireless IDS has to be different from traditional wired IDS. The main reason for this is explained below.

All wireless networks currently use the 802.11 IEEE standard of communication. So, when a client tries to communicate with an access point, both share the same frequency from a particular band in the range allowed by the 802.11 standard. They are also modulated by the same carrier

for proper communication. The 802.11 standard makes use of radio waves for modulation and communication.

Currently, wireless networks can operate in three main frequency bands: 900 MHz, 2.4 GHz, and 5.8 GHz. If you consider conventional FM radio, it is divided into a number of channels. We tune to specific channels like 98.3 and 94.1 MHz. Similarly, a wireless network uses a type of modulation known as FHSS (Frequency Hopping Spread Spectrum). The 2.4 GHz range is broken down into 79 different channels. It also uses the DSSS (Direct Sequence Spread Spectrum) type of modulation. Here, the 2.4 GHz range is broken into 11 channels. The available range is simply sub-divided into smaller parts to accommodate more channels or more simultaneous data transfer socket connections.

The reason I've made readers slog through the dreary math in the last paragraph is only to tell you that there is no single specific frequency in use at a given time. An IDS has to scan all the available channels in all the frequency bands available to sense the channels communicating at that time. So, even though the frequency range is standardised by the 802.11 standard, the detection is still very difficult. Therefore, the IDS must be able to scan all the channels, and that too, at a rapid rate so as to avoid loss of information regarding the intrusion.

We can configure an IDS to scan frequencies in only the band we use—802.11 a, b, g or n. To achieve this, the IDS uses a hopping system. (This is similar to frequency hopping in radio broadband communications and is beyond the scope of this article.) So, even though an IDS cannot collect all the data from all frequencies, it can only listen for the data from a channel when the frequency matches with that channel. During the remaining time, it cannot listen from that particular channel. After scanning all the remaining channels, it comes back to the initial one. (This is similar to time division multiplexing used in radio communications, again a subject that's beyond the scope of this article.)

Tracking rogue access points using Kismet

Irrespective of what kind of intrusion is taking place in a network, if the source of the problem can be found, then one can significantly eliminate the danger. A rogue AP (access point) is a Wi-Fi access point that is set up by an attacker. The purpose of a rogue AP is to sniff wireless network traffic in an effort to gain unauthorised access to the network environment under consideration. Once the attackers are in the network, they can disable services causing DoS attacks, collect sensitive information, or intercept secure data to be decrypted at their leisure.

There are many tools available to detect a rogue AP under Linux. A few of them are: Snort, Kismet and wIDS. Let's have an overview of rogue AP tracking with Kismet.

Kismet has all of the features that a normal packet filter has, but it also has many features dedicated to wireless networks. It has a built-in mechanism designed to detect any machine that's running NetStumbler, a program to gain access into a network. The software is also designed to decode WEP packets on the fly, as they are being captured. Crackers,

as we saw earlier, may spoof IP, MAC, and SSID. Kismet can fight back against these techniques because it supports SSID deauthentication and MAC spoofing.

Kismet is a fully capable and complete tool for detecting rogue APs on a network. It can be installed on a laptop, PC or a hand-held device with a wireless NIC (network interface card) to scan wireless packet data transfers. The tool comes with an added advantage: It can simultaneously access more than one NIC to collect packets. Such multiplexing is very helpful for collecting data at a rapid rate. As explained earlier, all the frequencies cannot be scanned simultaneously. However, by using multiple NICs one can speed up the process. Current versions of Kismet are limited to detecting only 802.11b wireless devices, and if you are in luck, Kismet might be able to detect an 802.11g device since it's backward-compatible with 802.11b.

As Kismet detects devices, it can predict the location of the rogue AP and plot it on a map if a GPS card is connected to the system. But this is experimental, as large metal objects and radio frequency interference attenuate the signal, making detection difficult. As wireless devices are being detected, the information to identify the node is recorded. For example, the SSID, the type of the device, manufacturer, etc, is recorded.

To actually find the location of the rogue AP, you must walk around in the network area with a wireless device that has Kismet, and let the tool detect any intrusion. Currently, having to move with the device is its only drawback. As you move closer and closer to the rogue point, the signal strength goes on increasing and the location can be pointed out.

Kismet is available as a free download from kismetwireless.net.

Tracking other types of intrusions using Kismet

Kismet can give detailed information about the types of attacks we have discussed, including DoS attacks or network breaches using tools such as NetStumbler. It also detects probing of networks by clients or nodes that are not part of the network, yet are sniffing the network. Kismet focuses on the 802.11 network standard. It also provides alerts to users when a problematic condition occurs.

Alerts provide a lot of help when someone is monitoring the network. Some are very useful and reliable, but others may generate false or useless information. The following is an example:

```
Alert name: NETSTUMBLER
Alert type: Fingerprint
Alert for: NetStumbler probe request
WVE: WVE-0000-0020
Alert message: NetStumbler (version) probe received from (IP address)
Tool-specific: yes (NetStumbler 1.22, 5:24, 3:30)
References: http://www.netstumbler.com
Details: In an attempt to sniff the SSID of a network
NetStumbler sends out probe packets. This is not good
in all situations, but when it is detected is potential
for take position is very low
```

What you can see above is a typical Kismet output (taken from the official Kismet documentation) when it finds some rogue client trying to access information from your network using a NetStumbler packet sniffer tool. To know all the types of intrusions Kismet can detect, please check the documentation at www.kismetwireless.net/documentation.shtml.


Wireless IDS drawbacks

The benefits of a wireless IDS are numerous, but there are a few drawbacks to be considered before deploying such a system in any network. Wireless intrusion detection is an experimental technology and proper evaluation is needed before its application. Because the technology is new, there are possibilities of bugs in the code, loopholes, or worse, vulnerabilities. These could potentially actually weaken the network if you fail to integrate the IDS effectively with your network. Having said that, with the current pace of research in this field, and newer and better algorithms and tools being released every day, most of the vulnerabilities will be gone in no time.

A wireless IDS is only as effective as the individuals who analyse and respond to the data gathered by the system (again, the network is only as strong as its weakest link/node). A wireless IDS, like a standard IDS, can require vast man, machine and money resources to analyse and respond to the collected data. Generally, more manpower is needed to monitor wireless IDS as opposed to standard IDS. The reason is simple: You not only have to analyse the data from various frequencies and bands, but also have to move around trying to find rogue APs. However, this point is debatable.

Given the advantages of these systems, these drawbacks (except the bugs, which could weaken the system) seem minor.

We have considered all the aspects in case of wireless IDS. Along with the strong encryption and the firewalls in networks, the IDS forms a second line of defence. However, wireless IDS, being a new technology, has its drawbacks, too. Then again, with the capability to detect probes, Denial of Service attacks, SSID spoof detection, and a variety of other 802.11 attacks, the benefits of a wireless IDS already outweigh the drawbacks.

Of course, IDS alone is not the solution to any security threat. It is a very small part of a much bigger security infrastructure. WLANs require a number of other security levels to be integrated before an adequate level of security can be achieved. With the current speed of adoption of new technologies, the ever increasing number of potential threats to a system, the growing complexity and sophistication of attacks, as well as attackers... you are doomed to lose your data unless you secure it using IDS, along with other lines of defence. 

By: Aditya Shevade

The author is a National Talent Scholar and a third year electronics engineering student, who takes a keen interest in programming and electronic design. A Linux user for the past two years, he enjoys playing the keyboard and is a good photographer. To know more about him, visit www.adityashevade.com.



Frederick Noronha

FOSS Campaigners Upbeat About Politicians' Stand on Software

Who is for FOSS, and who isn't? Either way, will pre-poll promises of 'digital inclusiveness' become a reality, post elections?

Indian prime minister-in-waiting and opposition leader, Lal Krishna Advani of the BJP seems to strongly favour Free Software and Open Source. While politicians are known to be lavish with pre-poll promises, Open Source campaigners were upbeat over the development.

New Delhi goes to the polls for its national parliamentary elections within weeks.

The Bharatiya Janata Party, a right-wing influential party that ruled India between 1998 to 2004, included "...key sections of the FOSS (Free/Libre and Open Source Software) Manifesto" in its IT vision that was released Saturday, March 14, 2009. Venkatesh Hariharan of Red Hat India and a FOSS lobbyist at the policy level, announced.

The BJP is centre-right in nature, and the party advocates conservative social policies, self-reliance, an emphasis on economic growth, with an emphasis on 'nationalism', and a strong military and defence policy.

Hariharan pointed to the just-released 40-page 'vision document' and said the man who could become India's next prime minister had 'strongly articulated' the party's views on FOSS in education, on open standards, on encouraging freely shareable FOSS-based knowledge and repositories like Wikipedia in Indian languages.

India's small but active FOSS campaigners have been pushing politicians and governments to take a stand in favour of Free and Open Source Software, and related issues.

In what it called a 'FOSS Manifesto', at a site called *public-software.in*, the campaigners called on Indian political parties "...to make FOSS usage and promotion a central part of the IT, e-government and education plans in their election manifestos."

They argued that FOSS "...is software that is liberally licensed to grant the right of users to study, change, and improve its design through the availability of its source code. The open, inclusive and participatory nature of FOSS is a natural fit for the vibrant traditions of Indian

democracy and its emphasis on sharing knowledge."


To bolster their case, FOSS evangelists said that since "...software is the foundation of the digital economy," India's IT infrastructure should be built on FOSS and "...not on closed, proprietary software systems that enforce restrictive licenses, limit the freedom of users and encourage monopolistic behaviour."

In the recent past, parties on the Left have also supported such FOSS initiatives. The CPI (M) government in Kerala is known for its computer education policy based on Free Software, while FOSS campaigners have also sought to build links between the Left and their campaigns, on ideological grounds.

"The Left, specially the Communist Party of India (Marxists), has also been supportive in taking a stand against software patents in India, and in favour of open standards," Hariharan says. "Of the major parties, only the (currently ruling) Congress is left to take a stand."

In a blog post, Hariharan commented, "As a long time supporter of free and open source software, I am delighted to see a major Indian political party [the BJP] endorse [the FOSS Manifesto]. However, I am even more delighted to see that this endorsement is rooted in a comprehensive vision for India's development."

The BJP's press release has been put up at www.lkadvani.in/eng/content/view/799/281

The FOSS Manifesto, which campaigners are placing before politicians prior to India's upcoming crucial parliamentary elections, is at www.public-software.in/FOSS-manifesto 

About the author:

Frederick Noronha (FN) is an independent journalist based in Goa. See his blog at fredericknoronha.wordpress.com and encounter his links on Twitter at twitter.com/fn

The Art of Guard

Part 1

Your first steps with SELinux.

From the 6th century BC, around the time the *Art of War* was written by Sun Tzu, till today, the basic tenets of warfare (attack/offence) and security (defence) remain the same. What has changed significantly from those days is the terrain. It has evolved from land and sea, to air and into cyberspace.

Securing cyberspace, especially in this connected world of the Internet, is of prime concern to governments, organisations and individuals. In a knowledge economy, the real wealth is in information and the bits of data that provide it.

Securing access to that data is of fundamental importance to the very existence of an organisation. CIOs and IT managers worldwide therefore look at security as one of their main areas of concern. IT security is a vast subject and covers many areas—the foremost being network security (firewalls), data security (encryption, backup, etc), computing security (restricting physical access, patching OS vulnerabilities, etc) and application security. Much attention has been paid to network security, data security and computing security, and thus there are a lot of products in the market competing with each other providing security solutions in these spaces.

But simply buying and implementing the security products is not enough. As Bruce Schneier states in his remarkable book, *Secrets and Lies: Digital Security in a Networked World*, the two fundamental principles about digital security are:

- Security is a chain; it's only as secure as the weakest link.
- Security is a process, not a product.

The application is a very important link in the IT security chain. It is the application that the end-user consumes and it is the application interface that is exposed to all users (even those without proper login credentials). These applications control the use of resources granted to them by the OS—memory, storage, networks, etc. Most applications are complex and run into tens of thousands of lines of code and can contain innumerable vulnerabilities.

Moreover, typically, multiple applications are executed on the same operating system instance. To give a simple example, a mail server running SMTP/POP/IMAP will, in

all probability, also be running a Web server. The Web server interface could be used for configuration or even for Web mail purposes.

Vulnerabilities in one application could bring down the entire system affecting other services and processes, or worse, lead an attacker to a storehouse of confidential information.

Why SELinux?

Application software will have its flaws and bugs, and will remain like that in the near future, depending on the nature

of the application and the constraints/awareness of the developers.

If multiple applications running on the same OS have to be secured, the OS has to play a crucial role in defining the confines of these applications. Security can only be achieved with better underlying operating system security that can isolate applications and files used by each, thus protecting the integrity of the entire system.

In most organisations that have implemented some form of security and have a security policy in place, the weakest link in the security chain are the systems administrators. They generally have access to most files in the system and can—by accident or design—perform operations on confidential files. Confidentiality of data in modern systems is another pressing requirement.

The NSA (National Security Agency) of the US, which originally developed SELinux, states: "The Security-enhanced Linux's new features are designed to enforce the separation of information, based on confidentiality and integrity requirements. They are designed for preventing processes from reading data and programs, tampering with data and programs, bypassing application security mechanisms, executing untrustworthy programs, or interfering with other processes in violation of the systems security policy. They also help to confine the potential damage that can be caused by malicious or flawed programs. They should also be useful for enabling a single system to be used by users with differing security authorisations to access multiple kinds of information with differing security requirements without compromising those security requirements."

In other words, SELinux does implement some kind of Access Control Mechanism to achieve the above.



Traditional UNIX-like operating systems have had an Access-Control Mechanism that has remained, and still remains, one of the strongest security features of this family of operating systems.

We are all aware of the 'rwx' bits set on files and folders along with some special permissions. These permissions, along with user identity (UID) and group identity (GID), form the basis of traditional access control. This access control prevents unauthorised access to files and processes. *chroot* jails further confine filesystem access to a running process.

Is there a flaw in this age-old time-tested Access Control Mechanism? If no, what is the need for SELinux?

The inherent flaw in the traditional permissions model is that of DISCRETION. The owner of a particular file, for example, can change the permissions of an object, at will. Just imagine the following scenario: */etc/passwd*, the file that contains the user database, has default permissions of 644 and the file is owned by the root user, who—by accident or design—assigns it 666 permissions (global read and write). Any user could write into the user database, thus changing UIDs, home directories, etc. Imagine the security breach if the root user were to do such a thing.

The traditional Access Control Mechanism of permissions on files and processes is thus discretionary—it can be changed at the discretion of the owner and is often termed as Discretionary Access Control (DAC). Herein lies an inherent security flaw.

By contrast, SELinux implements Mandatory Access Control (MAC), where access control decisions are not at the discretion of individual users or even systems administrators.

Understanding MAC

Mandatory Access Control (MAC) ensures that the security policy of an organisation is adhered to. The various 'actors' in any security policy can broadly be classified as:

- Subjects
- Objects

Subjects perform certain operations on objects and the security policy specifies which of these operations are allowed or disallowed through Mandatory Access Control. In an operating system:

- Subjects are typically processes or threads
- Objects are typically files, directories, TCP/UDP ports, shared memory segments, network interfaces, etc

When an application process or thread tries to access objects that are disallowed in the security policy, an ACCESS DENIAL occurs, thus providing a layer of security between applications. Each application can be understood to be running in its domain, and subjects can typically access objects or perform operations within this domain, thereby confining the application within the boundaries of its 'domain'.

A new process/thread, on execution, might create its own domain through something referred to as domain transition. We will come to this in the later part of this series.

To identify subjects and objects and to determine what

access is allowed to whom, SELinux implements a MAC mechanism called Type Enforcement. This enables a granular control of the access mechanism over the operating system.

Type Enforcement and Security Contexts

Under Type Enforcement, certain attributes are applied to all objects and subjects. These attributes are termed as Security Contexts. Each process and file/directory/port on the system is assigned a Security Context based on which the Type Enforcement policy allows/disallows access.

Security Contexts are stored in Extended Attributes (xattrs) on an ext2/ext3 filesystem. A typical SELinux security context is of the form: *User Identity:Role:Type/Domain*

In multi-layer security and multi-category systems, two more attributes—sensitivity and category—are added. We will come to these later in the series.

Roles and users in the security context are present to support RBAC (Role-Based Access Control) features in SELinux. RBAC is to enable user privileges within the SELinux system. This too will be discussed later.

In this introductory article, our focus will be on SELinux Type Enforcement to give you a feel of SELinux, and thus we will primarily highlight the third attribute of the Security Context—Type/Domain.

Introduction to SELinux policies

The SELinux Type Enforcement Policy is based on rules that apply on Security Contexts.

Processes running with a certain Security Context are either allowed or disallowed to perform operations (access permissions—read/write/getattribute/setattribute, etc) on objects with certain Security Contexts.

The total compilation of these rules is called a SELinux policy. Red Hat Enterprise Linux comes with two standard policies—Targeted and Strict.

To briefly summarise the above:

- All processes (subjects), files, TCP/UDP ports, etc, (objects) are assigned a security context.
- Rules are created that allow/disallow access type (read/write, etc) by subjects to objects based on Security Contexts of the subject and that of the object.
- The set of these rules creates a policy. Currently, we are discussing Type Enforcement Policies.
- This policy is applied mandatorily to the system and cannot be overridden by users and even system administrators.
- Mandatory Access Control through Type Enforcement Policies enables SELinux fine-grained control over the operating system and any applications that might be running on top of it—adding an extra defence layer.

Enabling SELinux

Before beginning to explore SELinux, we need to enable it in the system. I use Red Hat Enterprise Linux 5 on my desktop, though the concepts and most of the commands would be the same on other flavours as well.



Do you have the **EDGE**?

Yes / No / Can't Decide

Then submit your best IT implementation project details which you have completed / deployed for your company / client in the year 2008. We'll let you know whether you have the EDGE or Not!

To submit your Nominations please visit www.networkcomputing.in
Or email: srikanthrp@ubmindia.com / jamsheedg@ubmindia.com

Produced by

First, check whether SELinux is enabled on your system—as root, run the following command:

```
[root@vbg ~]# cat status
```

If your output is:

```
SELinux status: disabled
```

...then you will have to enable SELinux on your server, as it is currently disabled.

Else, SELinux should be enabled in your system and the output will state so. We will explore the meaning of this output in the next part of this series.

To enable SELinux, edit the `/etc/sysconfig/selinux` file. Initiate SELinux in 'permissive' mode:

```
SELINUX=permissive
```

And use the default targeted policy for SELinux Rules:

```
SELINUXTYPE=targeted
```

Once you set these values, save the file and reboot your system.

File relabelling

Once you have enabled SELinux in your system, you will notice that on the next reboot, all files in your system will be labelled. Labelling for files is akin to applying Extended Attributes required by SELinux to your ext2/ext3 filesystem. This process may take some time, so be patient and let the complete reboot occur.

Checking Security Contexts

Upon reboot, log in to your shell and issue the 'ls' command. Did you see any extended attributes? I doubt it.

To see the Security Context of files and directories, use `ls -Z`. You will see an output similar to the one below, containing the extended attributes:

```
[root@vbg ~]# ls -Z
-rw-rw-r-- 1 root root user_u:object_r:user_home_t:0 scratch-pad
-rw-rw-r-- 1 root root user_u:object_r:user_home_t:0 wwws-definition
-rw-rw-r-- 1 root root user_u:object_r:user_home_t:0 personal-articles
-rw-rw-r-- 1 root root user_u:object_r:user_home_t:0 The Art of Queue vbg
```

You can now see the Security Context of all files/folders in your system. In the above snippet, for the file `scratch-pad`, the DAC related attributes are: `owner(vbg)`, `group(vbg)`, and permissions (`664`); while the MAC related Extended Attributes (`xattrs`) that define the Security Context are: `user_u:object_r:user_home_t:0`. The fourth attribute in the Security Context defines the sensitivity. We will discuss this fourth attribute in the later part of the series.

To check the ID of the logged in user, issue the

command `id -Z`:

```
[root@vbg ~]# id -Z
root_u:system_u:unconfined_t:0
```

The `id` command also gives the UID/GID of the user:

```
[root@vbg ~]# id
uid=0(root) gid=0(root) groups=0(root) context=unconfined_t:0
```

To see a list of processes (subjects) with their security contexts, issue the `ps auxZ` command:


```
[root@vbg ~]# ps auxZ
user_u:system_u:unconfined_t:0 3543 ? Ss 0:01 /usr/bin/podman /usr/bin/podman --socketpath /dev/shm/podman.sock
user_u:system_u:unconfined_t:0 3546 ? S 0:02 /usr/bin/podman
user_u:system_u:unconfined_t:0 3750 pts/0 Ss 0:00 /usr/bin/podman
user_u:system_u:unconfined_t:0 3760 ? S 0:00 /usr/bin/podman
user_u:system_u:unconfined_t:0 3762 pts/0 Ss 0:02 /usr/bin/podman
root_u:system_u:httpd_t:0 root 4:27:41.0 1 4 4:27:46 2021A ps/2
R+ 18:46 1001 /usr/sbin/pdmpd
```

Please note that each process has been assigned a 'security_context'. Do also note the type assigned to these subjects.

Try out the above introductory commands in SELinux. Explore... there is a wealth of information available on SELinux on the Web. Do not enforce a Strict Policy while you are not too sure about SELinux Policies, which we will cover in detail in the coming issues. If you end up in a state that your system will not boot satisfactorily or if you would simply like to disable SELinux at boot time, press `e` on your GRUB screen to edit the boot sequence when rebooting. Edit the default kernel booting parameters by appending `"selinux=0"` at the end of the line and continue to boot.

This will temporarily disable SELinux and you will be back to traditional DAC security.

Coming up

- Enforcing and Permissive Modes
- Controlling SELinux
- Targeted Policy
- Understanding the Targeted Policy
- Policy Modules 

By: Varad Gupta

Varad is an open source enthusiast who strongly believes in the open source collaborative model not only for technology but also for business. India's first RHCSS (Red Hat Certified Security Specialist), he has been involved in spreading open source through Keen & Able Computers Pvt Ltd, an open source systems integration company, and FOSTERING Linux, a FOSS training, education and research training centre.

LINUX JOBS

Post: Solution Architect - Java
Company: Cybernet - StaffSupport
Profile: Should possess knowledge of Java/ J2EE, Servlets, JSP, EJB etc.
Exp: 8-12
Location: Chennai
Email: aarthimajid@cybernet.com

Post: Principal Engineer
Company: Yahoo Software Development India Pvt. Ltd.
Profile: Incumbent should possess 7+ yrs of experience in the same domain.
Exp: 7-12
Location: Bangalore
Email: logeshh@yahoo-inc.com

Post: Sr PHP Developer
Company: Nirvana Infosys India Pvt. Ltd.
Profile: Ideal candidate should have 2+ yrs of experience in the same.
Exp: 2-7
Location: Delhi/Gurgaon
Email: hr@nirvanainfo.com

Post: Component Developer
Company: Photon Infotech Pvt. Ltd.
Profile: Incumbent should have knowledge of J2EE, Net, Linux, Unix, Windows 2003 server etc.
Exp: 6-11
Location: Bangalore
Email: vmsuthkumar.s@photoninfotech.net

Post: Unix Administration
Company: Maples ESM Technologies Ltd.
Profile: Should have understanding of a Unix based operating system, preferably HP and Linux etc.
Exp: 2-5, **Location:** Pune
Email: pufair/aspt@maplesesm.com

Post: Software Programmer
Company: Customer is King
Profile: Should have a flair for managing the entire management.
Exp: 0-2, **Location:** Chennai
Email: customerisking@gmail.com

Post: PHP Developer
Company: HCL Technologies Ltd.
Profile: Applicant should have good knowledge of MySQL, XML, XHTML, with experience in Unix/Linux/FreeBSD etc.
Exp: 4-7
Location: Bangalore
Email: vinod.sylman@hcl.in

Post: Linux Device Drivers- M/ VSI
Company: Aspire Communications
Profile: Must be a engineering graduate with good communication skills.
Exp: 2-6, **Location:** Mysore
Email: hr-manager@aspirecom.net

Post: Oracle DBA
Company: Parth Infotech Ltd.
Profile: Should have working experience with Oracle 9i/ 10g & with ADO Unix/ Linux/ HP/UX/ Solaris.
Exp: 4-8
Location: Mumbai
Email: career.in@parthinfotech.com

Post: IT Manager
Company: Fidelity Information Services
Profile: Candidate should possess 10+ yrs of IT experience.
Exp: 10-17
Location: Gurgaon
Email: Deepak.bansal@fis.com

Post: Siebel Applications Analyst
Company: Fintax Corp.
Profile: BS in Computer Science or equivalent with 3+ yrs of experience in implementing/ supporting Siebel CRM and its customer modules etc.
Exp: 4-9
Location: Hyderabad
Email: anantha.gangai@fintax.com

Post: Embedded Software Engineer
Company: Carver People
Profile: BE Computer Science/ Electronics & Communications/ MCA.
Exp: 0-8
Location: Bangalore
Email: vamsi2008@gmail.com

Post: Unix Engineer
Company: Emeric Technologies Pvt. Ltd.
Profile: Ideal candidate should have 1+ yrs of experience in Unix Server systems administration with good knowledge of file system management, JVM, NFS, Patch installation etc.
Exp: 1-3
Location: Bangalore
Email: careers.in@emeric.com

Post: Developer
Company: Customer is King
Profile: Strong understanding and hands on experience in ASP.NET is a MUST
Exp: 0-3
Location: Chennai
Email: customerisking@gmail.com

Post: Sr. QA Engineer
Company: Oracle India Pvt. Ltd.
Profile: Bachelor degree in Computer Sciences or in related fields with experience in the Unix/ DOS/ Windows (50/50) platforms etc.
Exp: 5-9
Location: Bangalore
Email: prashant.k.singh@oracle.com

Post: Application Engineer
Company: Triazz Inc.
Profile: Candidates should have experience in architecting, developing and testing of modules on Linux/Unix. Experience with Standard Template Libraries, GTK is required.
Exp: 2-6
Location: Pune
Email: aaron@triazz.com

Post: Sun Solaris Administrator
Company: Laxmi & Suresh Infotech Ltd.
Profile: Should have 5+ yrs of experience in the same.
Exp: 3-10, **Location:** Mumbai
Email: lms.vashi@laxmiinfotech.com

Post: Developer MySQL
Company: First Indian Corporation Pvt. Ltd.
Profile: Incumbent should have experience with MySQL and DB2 on a Linux based platform. SQLServer 2005 etc.
Exp: 2-5, **Location:** Bangalore
Email: resumesindia@firstind.com

Post: PHP Programmer
Company: Corigate Infotech Ltd.
Profile: Should have advanced level knowledge of PHP/MySQL including e-commerce, Payment Gateways and PHP frameworks etc.
Exp: 2-3, **Location:** Delhi
Email: tangena@infotech.in

Post: Sr. Software Engineer
Company: Colaksoft Software Pvt. Ltd.
Profile: Should have a Bachelor's degree in Engineering or MCA with 6+ years of development experience in building enterprise grade applications/product.
Exp: 6-10, **Location:** Chennai
Email: mowbrend@colaksoft.net

Post: Unix Admin
Company: SAIC India
Profile: Strong background in UNIX/Linux with knowledge of Solaris is a must.
Exp: 3-8
Location: Noida
Email: jayd4.afra@saic.com

Post: System Deployment Engineer
Company: Task Hub Softtech Pvt. Ltd.
Profile: B.E (CS/ ECE), B.Tech, MCA and M.Sc with knowledge of Linux/ Unix.
Exp: 1-1, **Location:** Bangalore
Email: careers_taskhub2@yahoo.com

Post: Petrochemical Data Manager
Company: SAIC India Pvt. Ltd.
Profile: Bachelor's degree in petroleum engineering or earth science.
Exp: 3-8
Location: Noida
Email: sandeep.singh.bird@saic.com

Post: Sr. Engineer VAS
Company: Borsai Networks India
Profile: Degree in Electronics/ Computer Science/ Telecommunications Engineering with experience in deployment and maintenance of Large SMSC and other value added platform in GSM Environment.
Exp: 2-7
Location: Delhi/NCR
Email: jobs@borsainetworks.com



Get the Best Paying Jobs @ naukri.com
2,00,000 Jobs to choose from!

15,000 Recruiters | Get jobs by email | Confidentiality Assured

Powered by
naukri.com
India's No.1 Job Site

© 2008 Naukri.com. All rights reserved. Naukri.com is a registered trademark of Naukri.com. All other trademarks are the property of their respective owners.



USB 3.0

Buckle Up for Super Speed

What does the next version of USB have for us?

When I started my career a decade ago, my development PC used to have 32 MB RAM and 500 MB of hard disk space. The current generation of users can't even think of working under such configurations (or should I call it 'constraints'?).

Over the years, memory has gotten cheaper and bigger (with sizes in gigabytes now), and with different use-case models. Some of the most prominent use-case models are Flash drives that replaced floppy drives, gadgets with large memory to support media files, or camcorders with more recording time. As a user, you might want to move around large files between your PC and these devices.

With USB becoming the leading interface for PC and peripheral connectivity, speed is the next bottleneck that needs to be eliminated. As discussed, memory is getting cheaper and larger, and users are increasingly using memory devices to copy gigabyte-sized files. This would take frustratingly long hours with the existing interfaces. Also, the time taken to copy data will decrease the battery performance.

To cater to such needs, the latest USB 3.0 solution comes with a maximum speed of gigabytes in seconds, with low power consumption. Linux, for its part as a desktop

operating system, will also need to support USB 3.0 in the near future.

Let us now explore USB 3.0 from an end user's point of view, in terms of features and what has "proven" itself in Linux. For more technical insights, refer to the USB 3.0 specifications at www.usb.org.

Super speed USB 3.0

The USB 2.0 specification released in 2000 introduced high-speed devices. With the introduction of USB 3.0 in 2008, a new 'super speed' connection is being introduced. It is backward compatible with USB 2.0 and supports data rates of 5 Gbps (60 MBps). The new specification introduces new connectors (Figures 1 and 2).

Some of the salient features of the new USB specification are:

- **Speed:** 10x performance when compared to USB 2.0
- **Power efficiency:** Both of the device and host
- **Backward compatible:** Class drivers continue to work; legacy devices can be plugged in.

Linux support

The promoters of USB 3.0 have demonstrated USB 3.0 on Linux. They have modified the USB core and also included




Figure 1: Standard connectors with added pins

the xHCI driver. xHCI is the new Host Controller specification developed to support the USB 3.0 transport. The new driver layers look like the diagram in Figure 3.

The code patch is yet to get integrated with the main-line kernel. This is because the xHCI specification has not yet been made public. You can read more on the Linux Super Speed demo in the references mentioned below.

What can end users expect?

The real question now is: as a user of this technology, what can you expect? Most of us use thumb drives and mobile gadgets on a daily basis. We use these to copy files to and from our PCs. These copy processes take more time. With USB 3.0, in theory, the copy process could get 10 times faster. This means that with USB 3.0, you could copy a file that took 30 minutes, within 3 minutes. It has been demonstrated that you get transfer speeds that are at least three times faster, including latencies (hardware/software).

To put it in more technical terms, in the SuperSpeed Conference held last November in San Jose, California, the promoters have demonstrated transfer speeds of 318 MBps on Windows and 233 MBps (typically 125 MBps) on Linux. Compared to this, the file transfer speed of the current USB 2.0 is around 60 MBps. So as an end user, you should buckle up for better speeds while using these gadgets, which are expected to hit the market in 2010. PCs that support these gadgets could also hit the market at the same time.

New technologies introduce more changes in our life. New trends like cheaper memory enabled us to see the entry of gadgets like thumb drives and high-capacity media players. High-speed Internet connections enabled many of us to experience new features like voice over Internet protocol and podcasts.




Figure 2: Connectors for hard-held devices

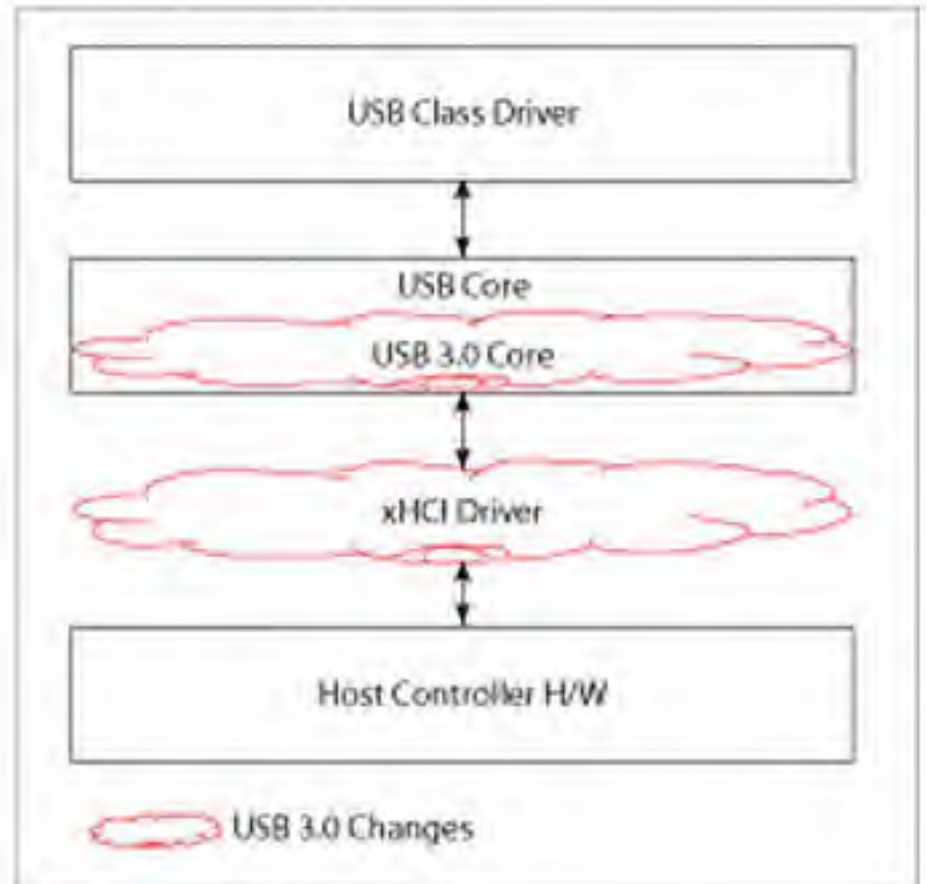



Figure 3: The new driver's layer in USB 3.0

As a popular peripheral connectivity solution, we hope higher speed USB 3.0 brings about better user experiences, apart from other benefits. 

References

- SuperSpeed USB Developers Conference Presentations: www.usb.org/developers/presentations
- USB 3.0 and Linux: sarah.thesharps.us/2008-12-07-13-35.cherry

By: [Rajaram Ragupathy](#)

The author welcomes your comments and feedback at rera_raja@yahoo.com

Got A LUG?

Visit www.openitis.com



Building IM Toys

Let's learn how to develop a GTalk bot to access digg.com stories. The bot is a hosted program that can be added as a friend in an IM client, and the user interacts with it through the IM window.

"The world loves toys; they are fun to fiddle with." —Swami Someoneand.

Add *digg@bot.im* as your GTalk friend and type 'help' in its IM window. I assure you that it's not me at the other end typing menu items and

with it through the IM window.

First let us define what our bot will do:

1. It will give the user a menu of numbered categories. The user needs to type the number of the category he/she is interested in.
2. The bot will take the user input and fetch the top 10 popular stories for that category. The story abstracts, with the corresponding links, will be displayed in the IM window.

Before we start, we will need:

1. A mechanism to make GTalk find and connect to our PHP script.
2. APIs to connect to Digg and fetch the stories.

imified [www.imified.com] is a service that facilitates the creation of IM bots. In short, it points a screen name on AIM, Yahoo, GTalk, and MSN chat services to our hosted script. You can register for a free account [here](http://www.imified.com).

Once you log in, go to *My Application*, then click on *Create Application*. Give your bot a GTalk Screen Name and a *General Identifier*. In the *Bot URL* field, enter the URL of your hosted PHP script. My script is hosted at <http://socialview.in/diggtoolkit/digg.php>. In the *Help Response* field, type in the response to be displayed when the user types 'help' in the IM window.

Download the PHP API from www.jaslabs.com/DownloadFiles/diggtoolkit.zip. Alternately, check <http://apidoc.digg.com/Toolkits> for APIs in a variety of languages. Go through the documentation to get familiar with the endpoints and the nomenclature. Thanks to its author, Justin Silverton, *diggclass.php* (present in *diggtoolkit.zip*) is well documented.

Our PHP script will have a simple logic:

1. Create an object of *diggclass*.

sending you Digg stories. A simple PHP script hosted on a Web server is accepting your messages and serving you content from digg.com while I attend classes, surf the Web and watch movies.

In this article, we'll learn how building this toy is both simple and fun.

An IM bot is a program that can be added as a friend in an Instant Messaging (IM) application, and the user interacts



2. Check the user input.
3. If the input is valid {
 - fetch stories
 - display stories }
- else {
 - show user the start menu }

And the script is as follows:

```
<?php
require_once 'diggsclass.php' // include the api
$diggapi = new diggclass(); // create the object
// we would send our bot to serving stories from these topics only
$topic = array( 1 => Technology, 2 => Lifestyle, 3 =>
Entertainment, 4 => Gaming, 5 => Sports, 6 => Science, 7 => World &
Business, 8 => Offbeat, 9 => );

$choice = $_REQUEST['req']; // get the user input
$choice = $choice + 0; // convert choice to a number

// if the choice was not between 1 and 9 show user the Menu
if($choice < 0 || $choice > 9)
    echo "Type in the topic number(1-9):";
    1 Technology (v) <br>
    2 Lifestyle (v) <br>
    3 Entertainment (v) <br>
    4 Gaming (v) <br>
    5 Sports (v) <br>
    6 Science (v) <br>
    7 World & Business (v) <br>
    8 Offbeat (v) <br>
    9 All (v) <br>
    return;

}

echo "Top 10 stories for choice { $choice } stories";
// fetch top stories 10 is default 3 returns popular stories
$stories = $diggapi->getStories(5 -> $topic[$choice]);
// iterate over each story
for($i=0;$i<count($stories);$i++)
    $desc = $stories[$i]['description'];
    $link = $stories[$i]['link'];
// display the description followed by the corresponding link
echo($i + 1). $desc. $link;
echo "\n\n";
}
```





Figure 1: The IM bot at work

>

Some interesting ideas that you can try for IM bots are listed below. You will have to hunt for the appropriate APIs or build the service yourself to implement these.


1. Dictionary and translation bot.
2. IMDB movie information bot—The user types in the movie name and the bot will return the cast, the synopsis, ratings and related information.
3. Movie showtime bot—The user can type in the city and the movie name, and the bot will return the show timings.
4. SMS bot—Typing on the tiny cell phone keypad is a pain. This bot can take in the text message with the recipient's name and deliver it.

So, aren't toys fun? 

By: Sagar Arlekar

The author is currently pursuing his MTech degree from IIT Bangalore. He works in the field of GIS and Web development. His interests are reading, cooking, sketching, travelling and admiring whatever good that exists around him. You can reach him at sagar.arlekar@gmail.com

Linuxforu.com
is making a come back soon.



'Good' Coding Practices

Here are some dos and don'ts while coding in C.

Following some sound logic is not always sufficient for developers working with C language to write good code.

There are some tacit assumptions when good programmers write programs. They typically make the code more reliable, portable, less error-prone and, most importantly, make it more readable. These standards vary depending on the project type, deployment scenario, compiler support, footprint limit, etc.

As far as the automotive industry is concerned, there exists an industry-wide solution called MISRA C standards, which contains a set of 141 rules that should be

followed when a piece of C code is written. (Read more about this at en.wikipedia.org/wiki/MISRA_C.)

These coding conventions are really good habits to follow, and they distinguish an amateur C programmer from a professional. In this article, we will try to understand such coding guidelines. The reference compiler here is GCC, but the guidelines are more or less the same for most of the other compilers.

General practices

Most projects have project specific guidelines regarding the variable naming and file naming conventions. The base

data types, such as 'int', 'char', 'unsigned short', etc, are not directly exposed to the programmers but they are 'defined' or 'typedefed' to `uint32_t`, `char8_t`, `uint16_t`, etc. The important storage class identifiers, such as 'static' and 'void', also are not directly exposed to the programmers but are #defined to 'VOID', 'STATIC', etc. The 'NULL', which also means 0x00, is indirectly used since one might need to point NULL to some place else, depending on the underlying architecture. All these 'indirections' are done to facilitate portability.

Status information, such as `STATUS_SUCCESS` (or the failures, if any), as well as the `TRUE` and `FALSE` macros are generally declared in a header file.

Mostly, one is expected to write only 80 characters per line of code and should have the tab space set to 4 characters. The editors are typically configured to 'auto expand' the tabs to spaces. In Vim editor, it could be done by *Esc +:set expandtab* or *Esc +:set noexpandtab*. The functions should not be lengthy and should avoid multiple return statements.

The file and function headers

Some projects have typical 'file header' and 'function header' formats. The following is an example:

```

/***** File Name *****/
** Includes
** Date created
** Author
** Contact
** Revision History
** Notes
*****/

/***** Source *****/

/* File:    InvolvePortFEM
 * (brief)  The function invokes the PortFEM and
 *          passes the output back to the caller.
 * (param)  The identifier of the current event.
 * (return)  STATUS_SUCCESS or STATUS_FAIL.
 * (note)   No memory allocation done in this function.
 */
STATUS InvolvePortFEM(IN karts_t current_event_id)

```

The aforementioned function header also caters to the 'doxygen' style of commenting, which helps in auto-generating the API documents. The fields that need to go in to these headers depend on the project requirements. Some projects also choose to put the copyright statements in the file, and it all changes from project to project.

Header file cautions

The header files should keep to the following format:

```

#ifndef _BASETYPES_H_
#define _BASETYPES_H_

```

```

/* File content goes here */

```

```

#ifdef _BASETYPES_H_

```

This avoids the compile time chaos caused due to multiple inclusions of the header files. If the header files are supposed to be written for a library that could be linked to C or C++ code, it should be written as follows:

```

#ifdef _BASETYPES_H_
#ifdef _BASETYPES_H_

#ifdef __cplusplus
extern "C" {
#endif /* __cplusplus */

/* File content goes here */

#ifdef __cplusplus
}
#endif /* __cplusplus */
#endif /* _BASETYPES_H_ */

```

Header file inclusions

The header files are often dependant upon each other. Hence, to avoid confusion, it is better to have a list of dependencies mentioned in the corresponding header files itself. Typically, people try to avoid including one header file into another directly, and let the source (C/C++) files include them in a correct sequence. Also, it is not advisable to #include one source file into another.

Source files should include header files in the following manner:

```

/***** Source *****/
#include <base size header files>
#include <dependant abc header files>
#include <base external/global specific header files>
#include <dependant external/project specific header files>

```

The sequence of header inclusions is extremely important and could save a lot of compilation time if done correctly.

No magic numbers

One more tacit rule in almost all typical projects is: 'No Magic Numbers'. Ideally, macros should be declared for all numeric constants that will be used in the program. This has to do with the scalability, and the same could cater to a scaled up or scaled down version of the product.

To give an example: today a code is written to scan through the 1k buffer. If it expands tomorrow, it could be easily made to run for the 2k buffer size or 512 bytes buffer size just by changing: *#define MAX_BUF_SIZE 1024* to *#define MAX_BUF_SIZE 2048* or *#define MAX_BUF_SIZE 512*.

Multilane macros

First of all, a judicious decision has to be taken when a multilane activity needs to be enclosed in a macro or in a function. If unsure, use the keyword `inline` and let the compiler make the decision.

Now, once it's been decided to write a multilane macro, it is preferable to enclose it in a `do{}while(0)`. There's no semicolon after it. The semicolon is supposed to be given at the time of macro call.

This ensures that we can convert the macro into a function at any point in time, and we need not browse through the entire code or keep adding semicolons at a later point of time.

For example:

```
#define INT(x, y) do{ x=0;
    int y=0;
}while(y!=1/while(0))
```

The reasons for this *do-while(0)* business are:

- 1) The macro becomes a single logical statement for the compiler, and hence can be safely invoked as:
- ```
if (condition)
 do_something();
else
 /* something else */
```
- 2) The macro is guaranteed to execute once.
  - 3) The compiler almost always optimises out the *do-while* and hence there's no performance overhead.
  - 4) Incomplete statements in the macro throw compilation errors.

## Putting the constant on the left when comparing with a constant

Consider the following piece of C code:

```
#define DELIMITING_CHARACTER 0

/* somewhere down ... */
char top_char;
while(1)
{
 top_char = getch();
 if(top_char == DELIMITING_CHARACTER)
 break;
}

printf("%c", top_char);
```

This works perfectly. Now, what happens if you or a colleague mistakenly takes out a `'=='` sign from the `if` statement in a hurry? The compiler will happily compile the code and the *while(1)* will keep

on executing endlessly.

A simple solution is to write the comparison in the following way:

```
top_char = getch();
if(DELIMITING_CHARACTER == top_char)
 break;
```

Now even if the `'=='` becomes `'='`, the compiler will straight away throw a compilation error and attract your attention. Hence, when comparing with a constant, the constant should always be placed to the left side of comparison.

## Writing explicit parenthesis

An amateur programmer often tries to sound smart by dropping the parenthesis just because C allows it. But consider the following piece of code:

```
if(condition_outer)
{
 if(condition_inner)
 statement_1;
}
else
 statement_2;
```

In C, every `'else'` branch belongs to the innermost possible `if` statement, which in this example is `statement_2`, and will be executed if, and only if, the `condition_inner` is false. This is often not what the programmer expected.

Hence, it is suggested that programmers should use the explicit parenthesis at all the places:

```
if(condition_outer)
{
 if(condition_inner)
 statement_1;
}
else
 statement_2;

...OR...

if(condition_outer)
{
 if(condition_inner)
 statement_1;
}
else
 statement_2;
```

...whichever logic is correct. The use of explicit parenthesis makes the code more readable and easy to understand.



## Print appropriate debug information

Most of the professional projects have macros like `TRACE` or `DEBUG_PRINT` defined globally. They typically provide multiple debug levels like `DEBUG_CRITICAL`, `DEBUG_MEDIUM`, and `DEBUG_ALL`. At compile time, one could choose to have the binary compiled with the appropriate debug level.

A small trick could be:

```
#define TRACE printf "TRACE: (%s) (%s) (%s)\n" \
 __FILE__ __FUNCTION__ __LINE__
```

Now `TRACE("Value is %d", val);` simply expands to:

```
printf("TRACE: (%s) (%s) (%s)\n" \
 __FILE__ __FUNCTION__ __LINE__, "Value is %d", val);
```

In addition to function and file information, one can choose to add line numbers and timestamps into the `TRACE` prints, depending upon the project requirements.


## Writing comments

Lastly, a very important point is to write precise comments. The prime motive behind writing comments is to make it easier for others to understand the code.

Some projects actually produce the API documentation from the code itself using tools like 'doxygen', and hence the commenting also needs to have specific formats.

In general, authors should keep in mind the target audience/maintainers when they write the comments, and hence the comments should focus more on the core algorithm/central logic than the language constructs. It is also very important to write the comments about any deviation/hacks made during implementation.

Apart from the source files, even the makefiles, build scripts, etc, that are used for a project, should also have sufficient comments and README files available in them.

Writing a proficient C code is not a skill to be learnt overnight but something we learn when we code and read the code written by the gurus. The source code of the Linux kernel is one such project where programmers get to learn a lot about smart coding practices. 

By: Nitesh Govande

The author is a Linux enthusiast and can be contacted at [niteshgovande@yahoo.com](mailto:niteshgovande@yahoo.com). His areas of interest include Linux systems software, application development and virtualisation. He is currently working with the LSI Research & Development Centre, Pune.

# PLEXUS

Hardware & Networking

Courses Offered

# RHCT RHCE



- ✦ Excellent Infrastructure.
- ✦ Certified Faculties.
- ✦ New batch starts on every Monday.
- ✦ Register for RHCE exam and avail **5 DAYS FREE** intensive examination coaching .

**Plexus Internetworking  
Technologies Pvt.Ltd.**

#31/14, Burkit Road, T-Nagar  
Chennai-600 017.

Ph:044-2433 7355 Mob: 99401 66510

Email: [info@plexus.net.in](mailto:info@plexus.net.in)

URL: [www.plexus.net.in](http://www.plexus.net.in)

Other Technical Partners





# You said it...



In the October 2008 issue of *LFY*, I got a 7-in-1 distro DVD. I would like to know how I can make the CD bootable or install the distros from the DVD on a hard drive.

—**Keyur Goswami, Jodhpur**

**ED:** The files in that DVD are ISO images. You need to burn each and every image to a CD/DVD in order to use them. Check your CD/DVD burning software for an option that lets you burn ISO images—you HAVE TO select this option, as simply burning one file on a CD won't work. We would recommend consulting the documentation of your CD burner application to learn how to burn ISO images, in case you haven't done it before.



I recently bought a laptop in which I had installed Fedora 10 from the DVD released with the January 09 issue of *LFY*. The OS is working fine, but when I wrote a C program using *vi* and tried to compile it using GCC, I got the following error message: *gcc: command not found*. I faced the same problem with a C++ program that I tried to compile with *g++*.

I am new to Linux. My only exposure to it has been to compile C and C++ programs using GCC and *g++*. I was under the impression that GCC and *g++* compilers are inbuilt in the Linux OS. I hope you can help me out in my problem. I am very interested in learning the internals of Linux, and without a good knowledge of its systems calls, I think it would be useless.

—**Sameer Sinha, by e-mail**

**Sankarshan Mukhopadhyay replies:** Would it be possible for you to check if you have GCC installed by using the following command as the root user from a console/terminal?

```
yum list | grep gcc
```

If you have it installed, you would

get an entry like what's shown below:

```
gcc386 4.3.2-7 installed
```

Else, choose *System* → *Administration* → *Add/Remove Software* and select the *Programming* group on the left hand side to select and install the packages.

You can also ask more questions on #fedora-india, on *irc.freenode.net* or the, *www.redhat.com/mailman/listinfo/fedora-india* mailing list. Happy hacking!



I was reading "The Open Movement..." by Dr M. Sasikumar [*LFY* March 2009, Page 40] and wanted to share some information with you. We have a school with about 1,000 kids in Jadan, Pali (Rajasthan) and I tried a lot to use open source programs in the school, but time and again I came up against the same problem. The curriculum in the school for computers is: MS Paint, MS Office, Foxpro and so on. How does one fight this?

I even installed a dual boot system to boot into Linux, set up a network in Linux, only to find teachers waiting for the boot screen to come up and then select Windows from it! When I removed that option, they stopped using that computer. There is sense in fighting this attitude, but only if there is something in the curriculum that encourages the use of Linux.

On the other hand, I tried *schooltool* as suggested in the article. I even managed to install the *Centre* program from *www.miller-group.net* which I would really recommend, but there is no fee module in it, and the school administration refused to use it and went ahead and bought software for Rs 1-2 lakhs.

So, it is not easy and reading your magazine is the only thing that kept me going to continue fighting such an attitude.

—**Gyaneshwar, by e-mail**

**ED:** Thank you for your valuable feedback about the drawbacks in our school curriculum. We can certainly relate to your frustration. The good news is that some political parties have specially included Open Source in their political manifestos for the 2009 national polls, and we're keeping our fingers crossed that others would do the same. This will be a start to making the curriculum vendor neutral at least.



I started subscribing to your magazine from November 2008. I am a newbie when it comes to Linux. However, your magazine is quite informative. I have one small request: if I could get the PDFs of your older issues it would be very helpful.

—**V Sundaresan, BTVL Chennai**

**ED:** We have been talking about launching a dedicated magazine website for quite some time now. The good news is, our in-house team is already testing various CMS alternatives, and I guess we will have something up by early June 2009. Then we'll certainly start uploading articles from our back issues as regular HTML pages so that everyone can access them.



First of all, thanks a lot for publishing my small note in the March issue. I had forgotten to mention that the Linux kernel compiling was not for me but for newbies, as I felt that most of the recent articles on Linux kernel would go over the heads of a new Linux user. I have been using Linux for over four years now and have a collection of around 200-250 Linux distros of various types and varieties. Thanks to *LFY*, I can boast of almost 75 that have come from this great magazine. Recently, I've started a blog at *ananthys4geeks.blogspot.com* and since I owe a lot of



# The Intricacies of the System V Semaphore

This article explains the semaphore related data structures and suggests suitable programs to implement the semaphore set.

**S**emaphore is a principal synchronisation mechanism for concurrent access of shared resources. Semaphores are classified as binary and counting semaphores. A binary semaphore is used to implement mutual exclusion and the counting semaphore is used when many shared resources are to be synchronised.

For example, if we want to book a train ticket, the train data is in the central database so that many users can access it at the same time. This is why we have many counters to book our tickets. But if two people are allowed to book a ticket to the same destination, on the same date and for the same train, then it could lead into problems if both receive a ticket with the same seat number. To avoid this, when the first user enters the records to book the ticket, the other is blocked until the first user exits. To accomplish this kind of environment, a binary semaphore is used.

Another approach of semaphores is resource counting. If there are many shared resources to be synchronised, then a counting semaphore is a right choice. Consider an example of a university that announces the results of graduating students on its website. If a count is set to

a specific number, say 1,000, this indicates that only 1,000 students can log into the website concurrently and check their results. Thereafter, any subsequent access will be locked till the count is less than a thousand.

In Linux, the System V Semaphore creates a semaphore set where many semaphores can be created, each being a binary or a counting semaphore, depending on the requirement.

Now let's consider creating a semaphore set with one binary and one counting semaphore. To understand this, let us consider another real world scenario: a team of 10 members working offshore have to take up a conference call from a client who is on-site. They all get connected to the conference call, so we need to have a counting semaphore that will have to maintain a count of the team members and synchronise access to common resources between them. But it has been noticed that during a conference call, only one person is given access to talk at a particular time and the rest are not allowed access. So a binary semaphore is required in order to achieve mutual exclusion amongst the team members.

Considering this scenario, this article explains the semaphore related data structures and suggests suitable programs to implement the semaphore set, which may be of the binary



or counting type (or a combination of the two), depending on the requirements.

In our previous article, "The Basics of the System V Semaphore", *LFY, July 2006*, page 86-89, we explained the binary semaphore's primitives; this article will explain the intricacies of complex semaphore programs.

Different functions are required to create and access a semaphore. To create a semaphore set, the *semget* function is used:

```
int semget (key_t key, int Number of Semaphores, int SemaphoreFlag)
```

The first argument *key* is the unique identifier for an IPC resource. The second argument is the number of semaphores in the semaphore set, and the third argument is the semaphore flag. The flag *IPC\_CREAT* is used to create a new semaphore set. When *IPC\_CREAT | IPC\_EXCL* is used, the system call returns -1 if the semaphore set already exists. Otherwise, the call creates a new semaphore set and returns zero.

*semctl()* is a system call used to perform control operations on the semaphore set. The syntax of the call is as follows:

```
int semctl (int SemaphoreID, int SemaphoreNumber, int Command, union semun)
```

...where the structure of *union semun* is:

```
union semun {
 int val;
 struct semid_ds *buffer;
 unsigned short *array;
}
```

The first member of the *union semun*, *val*, sets the value of a semaphore. Here is where programmers decide whether they need the binary or counting semaphore. If the *val* is assigned to 1, then it becomes a binary semaphore and is used to perform mutual exclusion; if the value is greater than 1 then the semaphore is of the counting type and the initial value of the first member is set to be equivalent to the number of shared resources. The second member, the *buffer*, is a data structure describing a set of semaphores. The next member is the *array*, which is used to set or get the value for all the created semaphores by passing the *SETALL* or *GETALL* command correspondingly to *semctl* function.

The following code snippet [Program 1] explains how to create two semaphores in the semaphore set, in which one is a binary and the other a counting semaphore with an initial value of 2.

```
1 /* Semidctl.c - This program demonstrates
2 * creation of multiple semaphores */
3 #include <stdio.h>
4 #include <sys/types.h>
5
```

```
6 #define SEMID1 1
7 #define SEMID2 2
8 #define SEMID3 3
9 #define SEMID4 4
10
11 #define SETALL 0
12 #define GETALL 1
13 #define SEMOP_ARRAY_SIZE 2
14
15 #define SEMOP_INCR 1
16 #define SEMOP_DECR 2
17 #define SEMOP_WAIT 3
18 #define SEMOP_NOWAIT 4
19 #define SEMOP_GETVAL 5
20 #define SEMOP_SETVAL 6
21 #define SEMOP_GETPID 7
22 #define SEMOP_SETPID 8
23
```

Notice that the semaphores have been created by passing the *SETALL* command as seen in Line 18 in the *semctl* function, and are retrieved through the *GETVAL* function in Line 20.

Now let us look at how to use these semaphores for synchronisation between two different processes. Locking or unlocking a critical section by means of decrementing or incrementing the semaphore's count value is done by the *semop* system call. The increment and the decrement operations are executed atomically on selected members of the semaphore set indicated by the *SemaphoreID*. The second argument of the *semop* system call is a pointer to the array of semaphore operations structure *sembuf*, which has the following three members:

```
struct sembuf {
 short sem_num;
 short sem_op;
 short sem_flg;
}
```

A set of operations takes place, each being performed on every semaphore in the semaphore set, where the first member of the *sembuf* structure is the semaphore number (here it is the first semaphore of the set, so it is *Semaphore 0*); and the second member, *sem\_op*, specifies whether the semaphore operation is incrementing or decrementing the value of the semaphore. The third member is the *sem\_flg*, which is the operation flag that is used to specify whether the calling process can wait for the semaphore or not.

The third argument *Nbuffer* of the *semop* system call specifies the number of *sembuf* structures in the array.

A particular process can use a binary semaphore to access a critical section, but what if the process encounters more than one critical section? In that case, there should be more than one semaphore to synchronise multiple critical sections. Linux provides



a semaphore set that can contain more than one semaphore to access multiple critical sections.

The program seen in the following code snippet [Program 2] demonstrates how the semaphore set can be used for multiple critical sections using two semaphores.

```
1 /* Semaphore.c: This program will demonstrate synchronizing
2 * multiple critical sections using multiple semaphores */
3
4 #include <stdio.h>
5 #include <semaphore.h>
6
7 main() {
8 int keyval=0;
9 key = fcntl(0, F_GETFL);
10 semid = semget(key, 2, 0);
11 struct semid *s[2]={0, SEM_UNDO, SEM_UNDO};
12
13 printf("locking semaphore1\n");
14 semop(semid, &s[0], 1);
15 printf("Inside critical section 1---semaphore1\n");
16 printf("ENTER to unlock semaphore1\n");
17 getch();
18 if(s[0].sem_val == 1)
19 semop(semid, &s[0], 1);
20 printf("semaphore1 unlocked\n");
21
22 printf("locking semaphore2\n");
23 semop(semid, &s[1], 1);
24 printf("Inside critical section 2---semaphore2\n");
25 printf("ENTER to unlock semaphore2\n");
26 getch();
27 if(s[1].sem_val == 1)
28 semop(semid, &s[1], 1);
29 printf("semaphore2 unlocked\n");
30 }
```

As the semaphore set has already been created with two semaphores in Program 1, in Line 10 of Program 2 – in the *semget* system call—0 is passed as the last argument since the semaphore has already been created.

In Line 11 the structure is initialised with values. For the first structure, 0 is passed to indicate the first semaphore, and 1 to indicate the second semaphore in the second structure. The second member indicates that the value of the semaphore is decrementing, and the third member determines what the *semop* function will do if the semaphore is busy:

- If the value is 0, it will wait until the semaphore is available.
- If the value is *IPC\_NOWAIT*, it will not wait for the semaphore, but returns with an error.

Sometimes, if a process locks the semaphore and exits abnormally without unlocking, then the other processes that are waiting for the semaphore have to wait forever (in the case of *sem\_flg=0*). The created semaphore set becomes unusable and the processes need to be killed

forcibly. To solve this issue, *SEM\_UNDO* can be OR'ed with 0. This *SEM\_UNDO* flag will automatically release the semaphore if a process exits without releasing it.

The locking and unlocking operation can be performed by the *semop* system call.

Lines 17-19 are the virtual critical section-1. Programmers can write the critical section of their code here. The library function *getchar()* waits for the *Enter* key, and if the same executable program is being executed by another terminal, the control stops at line 15, as the critical section-1 is already locked. The second terminal user cannot see the output of Line 17. After the critical section-1 access is over, the value of the semaphore must be incremented. First the value of the semaphore is reassigned in the *sembuf* structure and then the *semop* system call is called, so that the next waiting process for the semaphore can enter the critical section after decrementing the value of the semaphore. Once the first process leaves the critical section-1, the first terminal user can see the output of Line 22, which indicates that the semaphore is released; so now the second terminal user enters the critical section and can see the output of Line 17.

Now that the binary semaphore is handled, the second semaphore, which is the counting type, should be considered. Accessing the second semaphore indicated as 1 in the *sembuf* structure, the process enters the critical section-2 after leaving the critical section-1, and the first terminal user can see the output of Line 26. The other terminal process that is accessing the 1st critical section then leaves it and enters the critical section-2. So the user in the other terminal also can see the output of Line 26. Now it can be seen that both the processes are able to enter the critical section-2, which is because a counting semaphore was created with a count of 2. If another terminal is opened and the same program is executed, and if the process tries to enter the critical section-2, then it will not be given access and at this terminal the user cannot see the output of Line 26 as it will be blocked from the second critical section.

Consider a scenario where two binary semaphores are created. This can be understood from the following code snippet [Program 3]:

```
1 /* SemaphoreLock.c: This program creates two binary semaphores */
2
3 #include <sys/types.h>
4 #include <unistd.h>
5
6 #include <sem.h>
7 #define VAL
8 #define semid_1 100
9 #define semid_2 101
10 #define 100
11
12 main() {
13 int keyval=0;
```



```

14 static const semaphore_t s1=SEM_1;
15
16 key=shmctl(0,0);
17 shared=shmget(key,2*IPC_CREAT*0744);
18 warray=shared;
19 semid=shmctl(shared,IPC_CREAT,0);
20 drive=shmctl(semid,IPC_CREAT,0);
21 nt=shmctl(semid,IPC_CREAT,0);
22 printf("pid=%d\n",getpid());
23 }
24 }

```

We'll now create Program 4 and 5 to create the 'dead lock' situation as can be understood from Table 1.

```

1 /*SemDeadLock.c--This program demonstrates the program
2 *sequence of locking leading to deadlock situation*/
3
4 #include <stdio.h>
5 #include <sys/sem.h>
6
7 main()
8 {
9 int key=shmctl(0,0);
10 shared=shmget(key,2*IPC_CREAT*0744);
11 printf("pid=%d\n",getpid());
12
13 printf("locking semaphore\n");
14 semop(shared,SEM_WAIT);
15 printf("inside critical section\n");
16 getch();
17
18 printf("unlocking semaphore\n");
19 semop(shared,SEM_UNDO);
20 printf("inside critical section\n");
21 getch();
22
23 if(semctl(shared,0,SEM_UNDO)=0)
24 semop(shared,SEM_UNDO);
25 printf("semaphore1 unlocked\n");
26
27 if(semctl(shared,1,SEM_UNDO)=0)
28 semop(shared,SEM_UNDO);
29 printf("semaphore2 unlocked\n");
30 }

```

```

1 /*SemDeadLock.c--This program demonstrates the program
2 *sequence of locking leading to deadlock situation*/
3
4 #include <stdio.h>
5 #include <sys/sem.h>
6
7 main()
8 {
9 int key=shmctl(0,0);

```


```

9 key=shmctl(0,0);
10 shared=shmget(key,2*IPC_CREAT*0744);
11 printf("pid=%d\n",getpid());
12
13 printf("locking semaphore\n");
14 semop(shared,SEM_WAIT);
15 printf("inside critical section\n");
16 getch();
17
18 printf("unlocking semaphore\n");
19 semop(shared,SEM_UNDO);
20 printf("inside critical section\n");
21 getch();
22
23 if(semctl(shared,0,SEM_UNDO)=0)
24 semop(shared,SEM_UNDO);
25 printf("semaphore2 unlocked\n");
26
27 if(semctl(shared,1,SEM_UNDO)=0)
28 semop(shared,SEM_UNDO);
29 printf("semaphore1 unlocked\n");
30 }

```

| Program 4   | Program 5   |
|-------------|-------------|
| Lock S1     | Lock S2     |
| getchar ( ) | getchar ( ) |
| Lock S2     | Lock S1     |
| getchar ( ) | getchar ( ) |
| Unlock S1   | Unlock S2   |
| Unlock S2   | Unlock S1   |

Table 1

**Implementation:** If the first process locks the first critical section using the first semaphore (S1) and later locks the second critical section using the second semaphore (S2), then unlock S1 and S2. The second process locks the first critical section using S2 and the second critical section is locked by S1. In this scenario, both the processes wait indefinitely. This leads to the classical race condition. This is depicted in Program 4 and Program 5. So, if you know how to create a 'dead lock' situation, then obviously you know how to avoid the unwanted dead lock. To avoid dead lock, don't use more than one lock inside a critical section. But if you want to do that, you should preserve the same sequence of locking in all the processes. 

#### Acknowledgement

The authors would like to thank Mr V Jayasurya for his critical reading of this manuscript.

#### By: Shobana V and Dr B Thangaraju

The authors work with Talent Transformation, Wipro Technologies, Bangalore. They can be reached at shobana.venkatesan@wipro.com and balet.reju@wipro.com, respectively.



Subversion (svn) is a popular free software version control system that is widely used by a lot of developers across the world. Many popular community sites like SourceForge.net, OpenOffice.org, Netbeans, Tigris, etc. provide Subversion hosting for collaborative development.

This article takes you through some of the important features of the Subversion 1.6 release that benefits both users and the developers.

### Tree conflicts

A tree conflict is visualised as a 'conflict on a directory', unlike a conflict on a file or property. Subversion 1.6 detects a tree conflict and intimates the user with the

help of *svn status*, where a tree conflict is shown by a "C" in the fourth column of the status output. A tree conflict must be manually resolved by the user, without which committing changes becomes impossible.

For example, two users, A and B, are at work on different working copies, *wcA* and *wcB*, respectively. User A makes some changes to *Foo.c* file and commits it to the repository. Simultaneously, User B moves the file *Foo.c* to *Bar.c*. Now, B tries to update her working copy, which will make *Foo.c* to go into a 'conflicted state'. In older releases of Subversion, this does not happen—*Foo.c* receives the changes done by User A which is scheduled for deletion and user is not aware of the tree conflict. In Subversion 1.6, the user must resolve the conflicted state of *Bar.c*, and

# A Walk Through Subversion

# 1.6

The latest Subversion was released recently, and it's time to find out what's new.





commit the changes to the repository after checking the changes from User A, which is illustrated in Figure 1.

There are many such scenarios that are handled in a better way in Subversion 1.6.

## Security improvements

In older Subversion clients, the passwords supplied by the user were cached in plain text under most of the operating systems, mainly \*NIXes, which posed a major security risk. With Subversion 1.6, the passwords are cached in an encrypted form in \*NIX. If you are a GNOME user and want to store Subversion passwords in encrypted form, then you can use the GNOME Keyring to do so, just as KDE users can make use of the KWallet utility. In order to use either of these, the Subversion binaries must be compiled with the option for GNOME Keyring or KWallet support.

When Subversion does not find any means (provided by the operating system) to store passwords in an encrypted form, it informs users that the passwords are going to be stored in plain text and prompts them to confirm whether they would like to

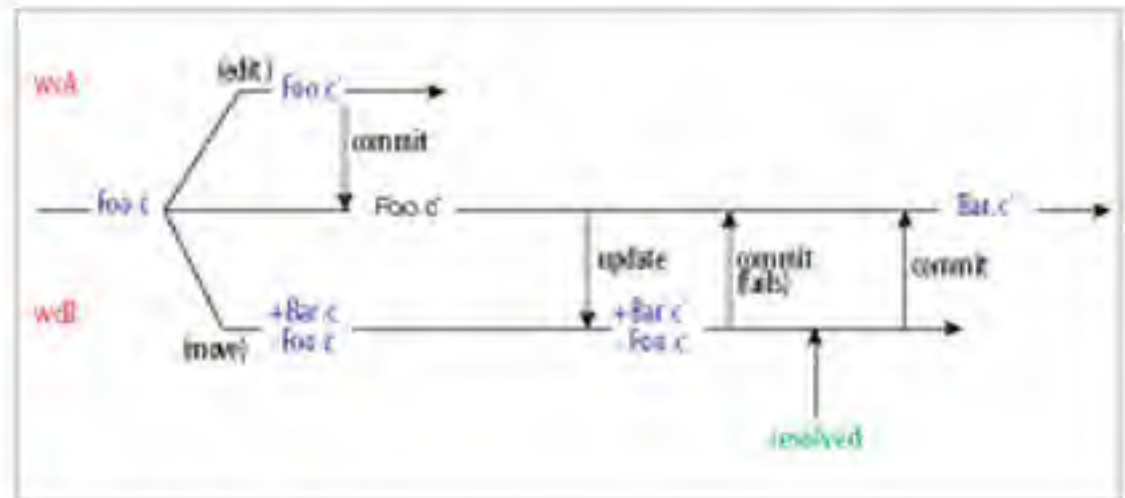


Figure 1: Tree conflict use case

proceed with this or not, as illustrated below:

\$ svn co http://localhost/svn/repos wc

Authentication realm = http://localhost:80: TEST SVN repository

Password for 'stylebot':

ATTENTION! Your password for authentication realm:

http://localhost:80: TEST SVN repository

can only be stored to disk unencrypted! You are advised to configure your system so that Subversion can store passwords encrypted, if



**Infrastructure Box**

For SMEs looking for a cost-effective but enterprise grade environment, we present the "Infrastructure Box". This comprises of Firewall, DNS, Proxy, Email, DHCP in a single physical box. This is a robust and cost-effective solution for you.

**Plugin Backup**

For those with an existing environment and requiring a Back-up Solution, look no further. Plug in our solution into your network and relax while it does your backup for you.

**Consulting**

We undertake Consulting for all IT/IS Infrastructure related requirements. Starting from the basic right to the advanced levels, we have the capability to deliver solutions that meet and exceed your expectations.

Please drop us an email at [enquiry@neuronit.in](mailto:enquiry@neuronit.in) to see how we can assist you further.

**Neuron IT Solutions**  
[www.neuronit.in](http://www.neuronit.in)  
 Building IT intelligently



**KONGU ENGINEERING COLLEGE**  
 PERUNDURAI, ERODE-638 052

(An autonomous Institution Affiliated to Anna University - Coimbatore)

**Ministry of Science and Technology,  
Govt. of India**

has sponsored us a study on

**"IT Professionals in  
Tamilnadu and Pondicherry".**

It involves preparing a directory of  
IT Professionals. You can be a part  
of this directory by filling in the details  
of the Proforma in our website

**www.itexpertsearch.in**

Prof. S.Kuppuswami  
B.E., M.Sc Engg., Dring(France)  
Principal

Thiruv. V.R. Sivasubramanian  
B.Com., B.L.,  
Correspondent



possible. See the documentation for details.

You can avoid future appearances of this warning by setting the value of the `--store-passwords` option to either `yes` or `no` in `$(HOME)/.subversion/servers`.

Store password `unencrypted(yes/no)? yes`

Checked out revision 0.

\$

This is really required because, in older releases of Subversion, the user is unaware of the plain text storage of passwords. The encrypted storage also applies for SSL client certificate passphrases in Subversion 1.6.

## Ctypes Python bindings

Subversion 1.6 has support for ctypes Python bindings, which have many advantages over the Swig-based Python bindings. The ctypes Python bindings are generated automatically, which provides access to all Subversion APIs and new features immediately, unlike Swig bindings which must be updated by the developer team for every release to get the latest APIs working.

The ctypes bindings are pure Python and cross-platform. You can generate the function definitions on one platform and copy them to another platform without incident. In other words, users don't need a compiler to install the binding. The ctypes bindings are both forward and backward compatible with different versions of Subversion and with different libraries, as long as the functions used in the programs have compatible definitions. Besides the low-level bindings, the ctypes bindings have high-level classes, which make it easy to access common Subversion functionality in a 'Pythonic' way.

## FSFS packing

Subversion 1.5 introduced the ability for FSFS repositories to be sharded into multiple directories. Subversion 1.6 takes the sharding concept further, and allows full shard directories to be *packed* into a single file. By reducing internal fragmentation in the filesystem, packed FSFS repositories have significant space savings over their unpacked counterparts, especially repositories that contain many small commits. Using a one-file-per-shard approach also allows Subversion to reduce disk I/O and better exploit operating system caches.

In order to pack a repository, we can use the `'svnadmin pack'` on the repository. Once the repository is packed, there is no means by which it can be brought back to the unpacked state, which prevents serving of this repository from servers older than Subversion 1.6.

## FS representation sharing

In a large repository with many branches and merges between these branches, it is common that many files will have similar history and content. In earlier versions of Subversion, this is stored as deltas against previous versions of the file, but Subversion 1.6 will use existing representation in the filesystem for duplicate storage. Though it is not an improvement visible to users, depending on the size of the repository and the degree of branching and merging, this can cause up to 20 and 15 per cent of space reduction for Berkeley DB and FSFS-based repositories, respectively.

## Single file externals

Subversion 1.6 supports single file externals. If the `svn:externals` URL points to a single file, this gets added to the working copy as a versioned item. The file externals cannot be moved or deleted like a normal versioned file, in which case the user must change the `svn:externals` property, but file externals could be copied. File externals can be identified by an 'x' in the switched status column of the `svn status` command.

## Command line syntax for relative URLs


Subversion 1.6 provides the user with a convenient way to refer the repository URLs, where a user need not type the entire URL of the repository, but can use a relative path inside the repository. The root of any Subversion repository can be defined as `"/`, the relative path inside the repository follows it. For example, if a user wants to refer a directory `sandbox` inside the trunk of the repository, the following notation could be used:

```
$ svn SUBCOMMIT -trunk/sandbox
```

## Other key features and improvements

In Subversion 1.6, there are many more new features available that are not covered in this article. They are:

- Logging support for `svnserve`
- BDB repository reverse deltas
- Support for memcache in FSFS repositories
- Apart from this, there are many more improvements to the existing features and lots of bug fixes.
- `svn log` can take multiple revisions
- The Trust server certificate option is added
- API changes and language bindings improvements

The latest release of Subversion (1.6) source can be downloaded from [subversion.tigris.org](http://subversion.tigris.org). If you want Subversion binaries for different platforms, visit [open.collab.net](http://open.collab.net). 

By: Senthil Kumaran S.

The author is part of CollabNet's Version Control group. He is a full committer of the Subversion project and a free software enthusiast. To know more, visit [www.stylesen.org](http://www.stylesen.org).





S.G. Ganesh

# Understanding 'Side-effects' in C/C++ Programs

Even experienced C/C++ developers don't entirely understand the side effects of operators. We cover this important topic this month.

I get mails on technical topics from readers with various backgrounds. A simple question from M Rajasekar prompted me to write about 'side-effects' in this column. He asked: "What is the output of the following program?"

```
int a=5;
printf("%d %d %d",a++,a++,a++);
```

As a student, I remember seeing such questions while writing written tests in campus placements (sometimes they are even asked in interviews). It just indicates how even experienced engineers who prepare the question paper, do not understand the question (nor do they know the answer!).

Here is a short answer to the question: the program has 'implementation defined behaviour'. The longer answer is: the expressions depend on 'side effects', and the resulting value of evaluating such expressions cannot be reliably predicted since the result can change with implementations.

So what are 'side-effects'? As the name indicates, they are other changes caused because of the evaluation of a main computation. In the statement "b = a++;", the main computation is the assignment of the value of *a* to *b*; the incremented value of *a* (because of *a++*) will be available only in the next statement and this is a 'side-effect'. If the statement is "a = a++;", it is a problem. Should the incremented value of *a* be (re)assigned to *a*, or should that result be lost? The C/C++ standards leave it to the implementation. Why?

To understand that, we need to know how an expression such as *a++* is implemented. While generating code, compilers put the value of *a* in a register and the temporary value of *a++* in another register. Once the current statement ends, the temporary value is assigned back to *a*. Now, for an expression such as "a = a++;", compilers can use different approaches. One compiler might use the register where the original value of *a* was stored; another compiler might use the register where the temporary value was stored. So, depending on the 'implementation' of the compiler, the result differs. C language gives this freedom to the compiler because the compiler can make use of this implementation choice to create the most efficient code for the given platform.

Let us look at an example to understand the unintuitive nature of programs that depend on side effects.


```
signed char a = 3;
while(a-- > 2);
printf("%d", a);
```

You can get the output as 43210, 4321 or even an infinite loop, printing 555555... depending on your 'implementation'!

To determine the impact of 'side effects' in an expression, C/C++ standards define 'sequence points'. Even experts find it difficult to understand the intricacies of sequence points, so we won't dare attempt to cover that topic here.

How about languages like Java and C#? As you might know, Java and C# do not have any 'implementation defined behaviour', which is good. For example, if the initial value of *a* is 2, after executing the statement "a = a++;", the value of *a* is 2; the program has a 'well-defined behaviour'.

Let us not debate whether the C/C++ approach or the Java/C# approach is better. It does not matter if the behaviour is implementation defined or well-defined—what is clear is that the behaviour is difficult to understand. If the initial value of *a* is 2, I initially expected that the value of *a* will be 3 after executing "a = a++;", in Java/C#. So, programs that depend on side effects are unintuitive and hard to understand. Don't ever write such code.

So what do you do if you're asked such a question in a written test or interview? Well, I don't know! Maybe send a mail to their HR team informing them that the question is 'silly, stupid and rubbish', though I guarantee that you won't get any response (and certainly not a job from them)! Better still, don't use their software—it's going to be as bad as their written test or interview questions! 

## About the author:

**S G Ganesh** is a research engineer in Siemens (Corporate Technology). His latest book is "60 Tips on Object Oriented Programming", published by Tata McGraw-Hill. You can reach him at [sgganesh@gmail.com](mailto:sgganesh@gmail.com).



## Programming in Python for Friends and Relations, Part 12

# Do You Speak my Language?



It is worth adding TTS to your applications, for no reason other than to attract attention.

"Grandpa, what was life like before computers got so small that they became invisible?" This caption of a cartoon sums up our predicament. How do we communicate with applications on our ever-shrinking computers? On the smaller machines, even familiar applications could look different and behave oddly. We may need help to use the different interface. However, reading a help document on a cramped Netbook or a smartphone screen is not easy. Wouldn't it be nicer if the 'mouse-over' text or a help message could be read out to us instead?

Efforts at text-to-speech have been around long before computers were created (see [en.wikipedia.org/wiki/Text\\_to\\_speech](http://en.wikipedia.org/wiki/Text_to_speech)). Getting the computer to speak in a nice voice has proven to be a remarkably hard

task. Many Interactive Voice Response (IVR) systems get around the problem by recording phrases spoken by people and playing back the appropriate sequence of wave files. This is clearly feasible if the size of the required vocabulary is small.

If you are willing to compromise and accept a voice that is comprehensible, even if robotic, there are a few options. The eSpeak ([espeak.sourceforge.net](http://espeak.sourceforge.net)) system should be the first to be explored, as it is widely used for its accessibility features and is the platform the OLPC/Sugar project uses. It is small and has support for a fair number of languages.

## Getting started

Python provides a wrapper for using the Speech Dispatcher server, which is generic for text to speech (TTS) applications. It can



talk to various TTS engines, including eSpeak, Flite and Festival. The last one has the best voices but these are easily available only for English.

You will need to ensure that several software packages are installed. The minimal list for Fedora 10 and Ubuntu 8.10 is:

- *espeak*
- *espeak-data* also on Ubuntu
- *speech-dispatcher*
- *speech-dispatcher-python* (on Fedora)
- *python-speechd* (on Ubuntu)

You may wish to enter some text in Hindi. Indic Onscreen Keyboard (available on Fedora, but not on Ubuntu) is a reasonable option. The on-screen keyboard and the keyboard layouts are provided by *lok* and *ml7n-contrib-hindi*.

You will need to configure the default engine in */etc/speech-dispatcher/speechd.conf* as follows:

```
DefaultModule espeak
```

Since Ubuntu and Fedora are now using PulseAudio, you should change the audio output in */etc/speech-dispatcher/modules/espeak.conf* to Pulse; otherwise, you may 'hear' only silence (on Fedora):

```
EspeakAudioOutputMethod "pulse"
```

You will need to create the default directory for the logs. The speech dispatcher will work but no log messages will be stored.

```
$ sudo mkdir /var/log/speech-dispatcher
```

Now, start the *speech-dispatcher* service and you are ready to start.

```
$ sudo service speech-dispatcher start
```

Verify that the speech dispatcher is working:

```
$ speak "Hello Hello Testing 1 2 3"
```

If you hear what you expect, you can now proceed further.

## Learning the first steps

You need to learn what the speech dispatcher provides and how it behaves. Interactive learning is the easiest. So, start your Python interpreter and try the following:

```
>>> import speechd
>>> start_speak()
>>> Hello(speechd.Speaker)
```

A little exploring and you realise that you need to

create an object of the type *Speaker*. So, create it and see what happens next:

```
>>> spk = speechd.Speaker(100)
>>> start_speak()
```

List commands are always helpful, so try the following:

```
>>> spk.list_output_modules()
['espeak -flow']
```

The first command tells you the output TTS engines available. If the default is not eSpeak or you wish to switch between the modules, you can easily do so:

```
>>> spk.set_output_module('espeak')
>>> spk.speak("Testing 1 2 3")
[225, OK MESSAGE QUEUED: (3, 0)]
```

That was easy, right? Now, list the voices available:

```
>>> spk.list_synthvoice_modules()
[('espeak', 'af-usmale'), ('espeak', 'af-usfemale'), ('espeak', 'af-usmale-25-year'), ('espeak', 'af-usfemale-25-year')]
>>> spk.set_synthvoice('af-usmale')
```

Hindi is an option, though only as a testing version, at present. Try to say something:

```
>>> spk.speak("Testing 1 2 3")
[225, OK MESSAGE QUEUED: (4, 0)]
```

You should have noticed a difference. Maybe it was too fast so let us slow it down:

```
>>> spk.set_rate(50)
>>> spk.speak("Testing 1 2 3")
```

Ascribing negative values while setting the rate, slows down the speech. The numbers should clearly be spoken in Hindi.

```
>>> spk.set_rate(-50)
>>> spk.set_voice('af-usfemale')
>>> spk.speak("Testing 1 2 3")
>>> spk.set_rate(150)
>>> spk.speak("Testing 1 2 3")
```

Continue your exploring. The voice isn't very feminine, so there's more potential for work!

What if you enter a Hindi text message? Try using the on-screen keyboard, *lok* and enter a Unicode message:

```
>>> spk.set_rate(100)
>>> spk.speak("नमस्ते")
```



# You said it...



my life's success to *LFY*. I have a dedicated link to the magazine.


Coming to the March issue, thanks for once again giving us a wonderful OS. I have been using Slackware since its 9.0 versions and to date, have never felt bored with this OS. I would especially like to thank Abhijit, who wrote a wonderful article on Slackware 12.2. It made me recall all the steps I went through. Kudos to him for not missing even a minor point in his article. Great work Abhijit!

Finally, another request – I have Debian 5 only. But would still like to have a copy of the full OS set from *LFY*. Hope you will not disappoint us!

—Ananth Gouri, by e-mail

**ED:** Wow! Thank you for all your compliments. We're glad that *LFY* has inspired you to keep abreast with Open Source over the years. In fact, your compliment inspires us to make the content more useful to address the requirements of the new generation of computer users—to get them to fall in love with free and open source software. By the way, Abhijit is one of our newest team members and the review was his first article—he walked around with a big grin after reading your feedback about the Slackware review!

Finally, we would have loved to bundle all the 5 Debian DVDs along with this issue; however budget constraints due to the current economic situation don't permit it. It's just DVD#1 with this issue, and we hope this will address the requirements of most of our readers, and the remaining few software packages (which we believe, shouldn't be too many) could be downloaded from the official Lenny repository.

 Hi, awesome team! Let me begin with the compliments first. There have been a lot of improvements in the articles over the

years. Give us more. I find veterans as well as students reading your magazine, and the number is increasing day by day. Kudos to your tech and marketing teams.

Now, for some points to improve upon. When you print article names on the cover page, please print the page numbers also so that we can jump to that page and start reading right away. No matter what article you publish, please insist on appending some links in the end, for further extensive reading. The matter you publish does not always satisfy an enthusiastic reader. Additionally, why haven't you published anything on XML and the semantic Web so far?

The March 09 edition carried a fatal typo. In the **Tips and Tricks** section (Page 96), under the "Lost Bash history" tip:

```
PROMPT_COMMAND='history <&'
```

They are backticks and not single quotes. I've noticed that your team has often confused the two. Some sound proof-checking should do the trick.

—Sabrahmanyam PVB  
(lamp19), systems engineer, Yahoo! Software (I) Ltd, Bangalore

**ED:** Thanks for your encouraging words—it feels great to hear that *LFY* is able to reach out to all sorts of audiences. All your suggestions seem very useful for our readers. We'll certainly try to implement these ASAP. As for articles on XML and the semantic Web, we haven't received any interesting contributions in these areas, of late. But we hope to make up for this in the forthcoming issues.

Thanks a tonne for pointing out the backticks error. Those should have been single quotes. It's an issue with the page-layout software we use for the design. However, yes, we ought to be more careful when we check the final proofs. We look forward to continued feedback from

you on areas we can improve on. Not only does it help us become more alert, it also helps our readers when we publish the corrections.

## Errata

In the February 2009 issue, the month was incorrectly spelled as 'Fabruary'. We regret the error.

In the March 2009 issue, in the Q&A section (page 17), in query number 3, the description of the cron entry was incorrect. We believe it happened while converting a diagram representation provided by the tech team into words by the edit team. Here's the correction: for a typical cron entry like the following,

```
* * * * * command
```

The starting five fields are as follows (from left to right):

1. Minutes (0-59)
2. Hours (0-23)
3. Day of month (1-31)
4. Month (1-12)
5. Day of Week (0-6), where 0=sunday

Thanks to Prabhat (prabhatrishi@rediffmail.com) for pointing out the mistake.

\*\*\*

In March 2009 Issue, Secure Communication article, page 84, column 2, the first code snippet had garbage text due to a printing issue. The correct code is as follows:

```
import sys
infile = open('salary_inp2.asp')
outfile = open('salary_sal_out.txt')
cra = sys.stdin.readlines()
sig = md5.decrypt_word(infile, outfile)
outfile.close()
infile.close()
```

Please send your comments or suggestions to:

**The Editor**

**LINUX FOR YOU Magazine**  
D-87/1, Okhla Industrial Area,  
Phase I, New Delhi 110020  
**Phone:** 011-26810601/02/03  
**Fax:** 011-26817553  
**Email:** lfedit@fyndia.com  
**Website:** www.openTjs.com



Don't be surprised if the Hindi text looks different in various consoles. By the way, try changing the TTS language to English and listen to the difference.

## Read aloud application

You can have the computer read aloud one of the first nursery rhymes that I learnt. Save it in a text file, *NurseryRhyme.txt*.

```

Auk, Auk, Auk, Auk
Auk, Auk, Auk, Auk
Auk, Auk, Auk, Auk
Auk, Auk, Auk, Auk

```

Your basic program, *read\_aloud.py*, would look like the following code:

```

import speech
s = speech.Speaker('ReadAloud')
s.set_output_unicode('espeak')
s.set_language('hi')
f = open('NurseryRhyme.txt')
s.set_rate(50)
for line in f.readlines():
 sentence = line.strip()
 print sentence
 s.speak(sentence)
s.close()

```

In Python 3, all strings will be Unicode, but currently, you will need to interpret the data read as a Unicode string.

You would notice that the entire poem is displayed even before the first line is spoken, but you want to display each line as it is being spoken. So, you need to think in terms of events. Your program needs to wait till the TTS engine has finished speaking the line. You would notice that the 'speak method' accepts a call back parameter and an events parameter. This ensures that the speech dispatcher will call your program back after the events you have requested.

The call-back method will be passed a parameter—the event, which resulted in it being called. The speech dispatcher currently has two events: *begin* and *end*. Needless to say, you would be interested only in the second event.

Python's threading module contains a class *Event* that will be very useful for this application. So the better version of your read aloud application will be what's shown below:

```

import speech
from threading import Event
def spoken(event):
 speech.close()
s = speech.Speaker('ReadAloud')

```

```

s.set_output_unicode('espeak')
s.set_language('hi')
f = open('NurseryRhyme.txt')
s.set_rate(50)
speech_over = Event()
for line in f.readlines():
 sentence = line.strip()
 print sentence
 s.speak(sentence, callback=spoken, event_type=END)
 speech_over.wait()
 speech_over.clear()
s.close()

```

You have added just six lines, modified one line and gained quite a bit of control over the ability to integrate the playing of sound with the rest of the application.

## Dhvani and Festival

Dhvani won the FOSS India Award last year and is available at [dhvani.sourceforge.net](http://dhvani.sourceforge.net). Using the source code from the repository and with a little help from Santhosh Thottingal, I could try the above examples of code. The voice quality seems more natural than the eSpeak voice. I hope the package evolves and becomes a part of the Fedora and Ubuntu repositories.

Information about integration with a speech dispatcher is available at [dhvani.sourceforge.net/doc/screenreader.html](http://dhvani.sourceforge.net/doc/screenreader.html). In our Python code, the change needed is:

```

s.set_output_unicode('dhvani-gemini')


```

Debian and Ubuntu have a package for a Hindi voice for Festival. The quality of voice is acceptable but eSpeak's is better.

If more applications used TTS, the speech quality would definitely improve. If the speech quality was better, more applications would have used TTS.

Fortunately, smart cell phones, Netbooks and ebook readers are going to change the dynamics of the current roadblock. Adding language support to the eSpeak project is likely to get acceptance the fastest. (On how to go about it, check out: [http://espeak.sourceforge.net/add\\_language.html](http://espeak.sourceforge.net/add_language.html))

If user applications rely on a speech dispatcher as Python encourages us to do, it is easier to use a TTS engine optimised for specific languages, e.g., Dhvani for Indic scripts. It is just a configuration detail.

It is worth adding TTS to your applications today, for no reason other than to attract attention. 

By: Dr. Anil Seth

The author is a consultant by profession and can be reached at [seth.anil@gmail.com](mailto:seth.anil@gmail.com)



**TECH@EDU EXPO**

EDUCATION. POWERED BY TECHNOLOGY.

13-14 June, 09, Pragati Maidan, New Delhi

Get a Glimpse into...

# The Future of Education

at Tech@Edu

**Mobiles for distance education. Digitisation of curriculum. Innovative teaching aids.****Video-conferencing powered classes. State-of-the-art audio & video equipment.****Labs that power innovation and R&D...**

...pretty exciting stuff—isn't it? But are these technologies really beneficial for educational institutions in India?

Have these been implemented by any of the institutions here? If yes, how's been the experience?

Most importantly, are these technologies worth the effort and investment?

These are some of the questions that will be addressed at Tech@Edu—a unique event aimed at showcasing the latest technologies that promise to benefit the Indian education sector. You will get to hear technology experts from the leading tech firms. Plus, there will be 'early adopters' from academia who will share their 'gyaan' based on their personal experiences.

## Workshops & Seminars

State-of-the-art LabsAudio & Video SystemsSimpler Management, Thanks to IT SolutionsHi-tech Teaching Aidse-Learning: Today & TomorrowHarnessing the Web and the MobileOpen Source Powered Education

## Organisations

- Schools
- Colleges
- Engineering institutions
- BBA/MBA colleges
- Education Trusts
- Education boards and societies
- Training institutions
- Examination bodies
- Government bodies

## Who should attend

- Chairmen & Directors
- Senior faculty
- Head of departments
- Professors
- Lecturers
- Teachers
- Purchase Heads
- Trustees
- Principals
- Deans

**Contact:**For speaker slots: Vandana Sharma - 9810295045  
email: [contact.techedu@efyindia.com](mailto:contact.techedu@efyindia.com)For Sponsorship: Devendra Kumar - 9810143735  
email: [sales.techedu@efyindia.com](mailto:sales.techedu@efyindia.com)**An  
EFY GROUP  
EVENT**[www.efyedu.com/techedu](http://www.efyedu.com/techedu)

EFY Enterprises Pvt Ltd, D-87/1, Okhla Industrial Area, Phase 1, New Delhi 110 020 Ph: 011-26810601-03; Fax: 011-26817563





Sandya Mannarswamy

Welcome to CodeSport. This month, we take a quick look at the problem of finding the strongly connected components of a directed graph. We then discuss the problem of finding the path-sum along the root-to-leaf path of a binary tree.

A huge 'thank you' to all the readers who sent in their comments to the problems we discussed in last month's column. We also had an open problem from March – on finding out whether a given graph with  $N$  nodes is connected, and to determine its best lower bound. The only question the algorithm can ask the adversary is of the form: "Does an edge exist between vertex  $u$  and vertex  $v$ ?" Readers were asked to come up with the best lower bound they could establish for this algorithm using an adversary argument. Congrats to Rajeev Kumar for sending in the correct answer to this problem.

Last month's takeaway problem was from graph theory. Given a directed graph, a strongly connected component (SCC) is a set of nodes such that every pair of nodes in the SCC is mutually reachable from each other. In simple terms, SCC corresponds to reducible loops in the graph. A directed graph can contain multiple strongly connected components. One of the interesting applications of SCC is to represent the loops in the software programs we write. For example, a *for* loop that you write in your C code is internally represented as a SCC when the compiler analyses the function for optimisation opportunities. Readers were asked to come up with an algorithm to find all the strongly connected components of the graph.

Since none of the solutions I received for the SCC problem were completely correct, I am going to keep this problem open for this month also. However, in order to help those who were pretty close to getting the correct answer, I will give you a clue. For most of the graph problems, one of the first places you should look is to see if you can use the 'depth first search' (DFS) traversal to solve it. Remember that DFS on a graph will produce a forest of trees. Let us recall the definition of SCC. A strongly connected component of a graph is a maximal set of nodes, along with their edges that are strongly connected. All nodes on a SCC are inter-reachable with each other.

## The SCC of a graph

If we represent each strongly connected component as a vertex, then the component graph has as its nodes each SCC in the original graph. An edge exists between two SCC nodes in the component graph if there is a directed edge from vertex  $u$  in one component to a vertex  $v$  in another component. Note that vertices  $u$  and  $v$  are NOT mutually reachable, for if there was a directed edge from  $u$  to  $v$  and from  $v$  to  $u$ , then they would have been selected as part of the same SCC. We can easily show that the component graph is a directed acyclic graph, and that each SCC component is maximal in the sense that no further vertices can be added to it. Since the SCC component graph is a DAG, it has a topological ordering and we can use it to enumerate the SCC components. Recall that we do topological sorting by listing the vertices with the longest DFS finishing time, first. Also note that if the vertices  $u$  and  $v$  are in two different SCCs, namely  $U$  and  $V$ , respectively, and if there is an edge from  $u$  to  $v$ , then  $F(U)$  is larger than  $F(V)$ , where  $F(U)$  is the largest finishing time for any vertex in  $U$ , and  $F(V)$  is the largest DFS finishing time for any vertex in  $V$ . This means that we can use topological ordering to output component  $V$  first and then component  $U$ .

## The transpose of a graph

Also recall that the transpose of a graph  $G$  is a graph containing the same set of vertices as in  $G$ , but an edge exists from vertex  $u$  to vertex  $v$  in the transpose graph  $G^T$ , if a directed edge exists from vertex  $v$  to vertex  $u$ . If we print the topological order of  $G$  in reverse order, then it is a topological order for  $G^T$ .

You should use the above facts to come up with the algorithm to find all the SCCs of graph  $G$ . Recall that since two vertices should be mutually reachable for them to be in a SCC, they will be present in the same tree  $T$  when we perform a 'depth first search' on either  $G$  or  $G^T$ . The strong clue is to do two DFSs, one on the original graph  $G$  and another on the transpose graph  $G^T$  and use



the information obtained from these to find the SCCs. These hints should be more than sufficient for readers to solve the problem. So do send in your solutions before the end of this month.

## This month's programming question

This month we will revisit binary trees. A binary tree is made of tree nodes, where each node contains a 'left' pointer, a 'right' pointer, and a data element. The 'root' pointer points to the topmost node in the tree. The left and right pointers themselves recursively point to smaller 'sub-trees' on either side. A null pointer represents a binary tree with no elements, namely, the empty tree.

You are given a binary tree containing a set of integer values, with no duplicates. There exist paths in the binary tree from the root to the leaves. Each path is distinguished by the set of nodes contained in the path. A root-to-leaf-path consists of a sequence of nodes in a tree, starting with the root node and proceeding downward to a leaf (a node with no children). An empty tree contains no root-to-leaf paths.

So, for example, the following tree has exactly four root-to-leaf paths:



Root-to-leaf paths:

- path 1: 5 4 11 7
- path 2: 5 4 11 2
- path 3: 5 8 13
- path 4: 5 8 4 1

We will consider the question of finding the sum of the values along each root-to-leaf path. For example, the sum of the values on the 5-4-11-7 path is  $5 + 4 + 11 + 7 = 27$ . We call the sum of values along the root-to-leaf path as the path-sum value. For instance, the path-sum values along the four paths in the binary tree above are 27, 22, 26 and 18.

Let us consider the problem associated with this, namely, given a sum value  $X$  and a binary tree, does there exist a root-to-leaf-path in the binary tree such that the path-sum value is equal to  $X$ ? For instance, if the sum value is 22, we can say that there exists no root-to-leaf-path in the above binary tree, which has a path-sum value equal to this. Our intention is to come up with an algorithm such that, given a binary tree and a sum, the algorithm returns true if the tree has a root-to-leaf path such that adding up all the values along the path equals the given sum. The algorithm returns false if no such path can be found.

The simplest approach seems to be the one where we traverse all the root-to-leaf paths, summing up the value along the path and checking when we reach a leaf, whether the sum matches the given value. A simple trick can reduce the processing time of the above algorithm. Instead of summing up the values from zero as we traverse each node along the root-to-leaf-path, we choose to subtract the value from the given sum value, as we traverse each node and pass the remaining sum value to the next node on the

path. If the remaining sum becomes zero or negative before we hit the leaf, we can terminate further traversal along that path and start exploring the next path. If the sum becomes zero, on traversing the leaf, we can return the answer 'true' since we know that the path-sum along the path matches the given sum value. Shown below is the pseudo code for this problem:

```

boolPathSumExists(pseudo_sum, struct node* tree)
{
 // return true if we have traversed past the leaf and the sum is equal to 0

 if (node == NULL)
 {
 return (sum == 0)
 }
 else
 {
 // otherwise check both subtrees
 if (sum == node->data)
 return (boolPathSumExists(sum, node->left)) ||
 (boolPathSumExists(sum, node->right))
 }
}


```

What is the complexity of the above algorithm? Is it possible to come up with a better algorithm? A variant of the path-sum problem is to enumerate all the paths from root-to-leaf in a binary tree. What is the worst-case time complexity of an algorithm that needs to enumerate all the root-to-leaf paths? I leave these questions for readers to think about.

## This month's takeaway problem

This month, let us stay with the binary trees. A 'binary search tree' (BST) or 'ordered binary tree' is a type of binary tree where the nodes are arranged in order: for each node in the tree, all elements in the left sub-tree of the node are less-or-equal to the value of the node, and all the elements in the right sub-tree of the node are greater than the node. Binary search trees are used for element insertion and look-up.

This month's problem is to determine whether a given binary tree is, in fact, a binary search tree or not. Given a plain binary tree, you have to examine the tree to determine if it meets the requirements of a binary search tree. To be a binary search tree, at every node, all of the nodes in its left tree must be  $\leq$  the node, and all of the nodes in its right subtree must be  $>$  the node. Your algorithm should return true if the binary tree meets the criteria for a binary search tree. Else it should return false.

If you have any favourite programming puzzles that you would like to discuss on this forum, please send them to me. Feel free to send your solutions and feedback to [sandyaasm\\_AT\\_yahoo\\_DOT\\_com](mailto:sandyaasm_AT_yahoo_DOT_com). Till we meet again next month, happy programming! 

### About the author:

**Sandya Mannerswamy** The author is a specialist in compiler optimisation and works at Hewlett-Packard India. She has a number of publications and patents to her credit, and her areas of interest include virtualisation technologies and software development tools.



# Getting Your Hands Dirty With Deb Files



As a Linux user, you have to deal with packages occasionally, if not every day. This article talks about the Debian way of going about it!

**F**irst things first: a package, put simply, is software created especially for a particular GNU/Linux distribution. Debian packages are called 'deb' in short, and have the *.deb* file extension.

There are many advantages of using packages, rather than the upstream sources, to install any software. Packages are always signed by its maintainers. Some security updates come faster in a distribution compared to its upstream, since it affects the whole distribution rather than a single package.

## Some internal details

Debian packages are handled by software called *dpkg*. Generally, users don't need to deal directly with *dpkg*. Its front-end is called *apt-get*, which takes care of almost all of your needs.

You can find all configuration files related to Apt package management in the */etc/apt/* directory. Let's look at each of the files:

- */etc/apt/sources.list* stores the list of repositories you are using. By default, this will be added, depending upon the Debian version you are using (stable, testing, unstable, etc). Make sure that you have security repositories enabled. You can comment, uncomment, add or delete any lines here, but don't forget to run *apt-get update* (see the package management section below) each time you edit this file.
- */etc/apt/apt.conf* is the general

configuration file where you can change many settings depending upon your need. If, for instance, you are behind a proxy, you can add the following line to your proxy settings:

```
Acquire::http::Proxy "http://username:password@IP:Port"
```

Please see the man page of *apt.conf* for a detailed description.

## Package management

Let's look at some practical examples. We know that *apt-get* is the front end to deal with Debian packages:

```
apt-get update # to update the external package database
apt-get install package_name
```

Now, if you are not sure about the exact *package\_name*, you can search in your package database:

```
apt-cache search package_name
```

You can also view more information about packages, using the following command:

```
apt-cache show package_name
```

This will display the package description, package size, maintainer information, etc, to help you decide if this is the package that you're looking for. Removing a package is also very easy:






Figure 1: Default UI of aptitude



Figure 2: The Synaptic package manager

```
apt-get remove package_name
```

By default, `apt-get` will not remove any configuration file for your package (to make the upgrade smooth). If you want to remove configuration files, issue the following command:

```
apt-get --purge remove package_name
```

### Update for good

I'm sure you would always want to keep your system updated just so that you're in sync with the latest features, as well as security updates:

```
apt-get update
```

```
apt-get upgrade # download & install the updates
```

If you want to upgrade your system to a newer Debian release (for example, from Etch to Lenny, or stable Lenny to unstable Sid), use the following command:

```
apt-get dist-upgrade
```

Of course, don't forget to do an `apt-get update` before the distribution upgrade.

### Advanced package management

Sometimes, you might want to play with packages, and want all build dependencies of packages in one shot:

```
apt-get build-dep package_name
```

You might also want to download the source of the Debian package:

```
apt-get source package_name
```

If you are a developer, and want to check out the development repository of a Debian package, you have a nice command called `debcheckout`. Make sure that you have the `devscripts` package installed for this.

**NACE™**  
move to the next level

TRAINING & CONSULTANCY

**BECOME AN  
OPEN SOURCE  
PROFESSIONAL.**

**IT INDUSTRY IS  
WAITING FOR...**



**LINUX  
APACHE  
MySQL  
PHP**

**Linux Basics  
Linux Administration  
Linux Networking  
Linux Security  
OpenOffice.org 3.0**



**CHENNAI**

**T.NAGAR  
99405 42222**

**ANNA NAGAR  
97909 25554**

**ADYAR  
9790925550**

**www.naceedu.com**



## Doing things graphically

GUI tools are more convenient for certain use cases. Debian includes many ways to manage packages. We will look at each of them. Let's start with the famous *aptitude*.

*aptitude* is a semi-graphical package management software, which comes with the power of the command line as well. To start *aptitude*, issue the following command:

```
aptitude
```

You'll notice something similar to the image in Figure 1. Although, it doesn't look very user friendly, it has lots of power. For instance, it can effectively scan your system and give you suggestions to remove some useless packages from your system.

Synaptic is a powerful graphical way to manage your packages. To start Synaptic, go to *System→Administration→Synaptic Package Manager*. Synaptic, as we all know, is easy to use. Pressing the 'reload' button will implement an *apt-get update*, while 'mark all updates' is the same as an *apt-get upgrade*. You can also manage offline installation with CDs/DVDs, using Synaptic. Go to *Edit→Add CD-ROM*. You can also manage the online software repositories with it. As you can see in Figure 2, it will even display a screenshot of the application with details on what the selected package does.

*Update-manager* is yet another graphical utility to manage easy package updates. After you log in to your desktop, by default, it runs in the background to check if your system packages are up to date or not. If there are updates, you get a GUI notification near your task bar.

## When things go wrong

We can't really say when things will go wrong, but they sometimes do, nonetheless. Apt has some options to correct corrupt installations. To get things in order, issue the following command:

```
apt-get -f install
```

Note that there is no argument after 'install'. It will automatically detect half-installed or broken packages and correct them.

You can find the status of all packages using Synaptic or by using the following command afterwards:

```
dpkg -f
```

The above command will give an output of the list of installed packages.

You can also view the list of files (including their locations) installed by a particular package, using the following command:

```
dpkg -f package_name
```

Just in case it's absolutely impossible to rectify a corrupt installed package, you can forcefully remove it from your system, using the code below:

```
dpkg --force-ai package_name
```

## "I want the latest software in Lenny!"

Sure. You can get it!

There are different ways to do so. The first option is to determine your requirements. Do you want only some software? Or do you want to use the latest software without going to 'unstable' or 'experimental' and take the risk of breaking your current stable working set-up?

Debian has a project called backports. This is recompiled software from 'testing' or 'unstable' branches for the 'stable' version. To get backports of software, simply add the following lines into your */etc/apt/sources.list* file:

```
deb http://www.backports.org/debian lenny-backports main contrib non-free
```

Run *apt-get update* and install a more recent version of your favourite application. For more information, visit *backports.org*.

Another way is to use 'pinning', which is a way to give higher priority to particular repositories for particular package(s). For example, if you have added the 'experimental' repository in your */etc/apt/sources.list* and want to install only 'mutt' from it, you can add the following lines into */etc/apt/preferences*; create this file, if it is not already there:

```
Package: mutt
```

```
Pin: release a=experimental
```

```
Pin-Priority: 999
```

Now, when you execute *apt-get install mutt* it will fetch 'mutt' from 'experimental'.

We have gone over some basics of Debian packages here. You can always visit *www.debian.org* for more details. There are some books available too, for example, *Debian 3.1 GNU/Linux Bible* by Jaldhar Vyas, or *The Debian System* by Martin Krafft. The best way to learn more about Debian is to install Debian and use it, practically :-)



By: Kartik Mistry

Kartik is a Debian developer and is deeply involved with various FOSS projects, particularly using i18n/l10n strings. You can contact him at [kartik@debian.org](mailto:kartik@debian.org) or check out his blog [fibs.wordpress.com](http://fibs.wordpress.com). He even has a Gujarati blog at [kartikm.wordpress.com](http://kartikm.wordpress.com).



# We have crossed **2.5 million** hits per month !



Join These Growing Numbers And Benefit...

## **electronicsforu.com**

**A comprehensive portal on electronics**

### Popular sections



#### **Circuit Lab**

Hundreds of circuit diagrams of different projects, for **FREE**



#### **Top 100**

Most popular technology websites, measured by their hits



#### **Events**

A comprehensive list of technical events, in India and abroad

Looking for something related to electronics? Try electronicsforu.com. This website has circuit diagrams, technical articles, bulletin board, dedicated index of electronics websites, list of popular websites ranked by hits...you name it, it's got it.

Check it out today!



EFY Enterprises Pvt Ltd D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020  
Phone: (011) 26810601-03 Fax: (011) 26817563/26812312 Email : info@efyindia.com



# A Practical Real-World Approach to Embedded Linux

The book definitely deserves shelf space in the library of anyone working in Embedded Linux.

In my humble opinion, *Embedded Linux Primer: A Practical Real-World Approach* is perhaps not the best place a newbie would want to get started, since it lacks a hands-on style to holistic embedded Linux development. The hands-on style is, however, found in places where the author has talked about BusyBox, some of the profiling and debugging tools, and the like. For anyone moderately familiar with the Linux kernel sources and concepts, a number of chapters should be easy browsing.

The book starts off with an introductory chapter and talks a bit about the state of embedded Linux, its soaring popularity and the reasons behind it.

The next chapter, titled 'Your First Embedded Experience', familiarises the reader with the basic anatomy of an embedded Linux system—the initialisation process, the storage options available, cross-development, and concludes by taking a look at some embedded Linux distributions.

Chapter 3 talks about 'Processor Basics'. As the name suggests, it familiarises the reader with the various processor architectures—from standalone processors and integrated processors, to the hardware platforms used in embedded systems.

The chapter titled 'The Linux Kernel—A Different Perspective', introduces the compilation of the kernel sources—the Linux kernel build system and configuration options are introduced here. There's nothing new for someone already familiar with how to compile the Linux kernel.

Chapter 5 (Kernel Initialisation) and Chapter 6 (System Initialisation),



together show how the Linux kernel and filesystems fit into the picture and give you a working environment. These chapters talk in some detail about the *init* process, initial RAM disk, and the Linux filesystems—Filesystem Hierarchy Standard (FHS), root filesystem, minimal filesystems, etc.

Chapter 7 is dedicated to bootloaders and their role in an embedded system. Together with Chapters 4 and 5, the reader gets a clear idea of how the various pieces fit in a working embedded Linux system. Chapter 8 talks about Linux device drivers.

The next one is dedicated to the popular filesystems being used in embedded Linux systems today. This chapter also covers JFSS2, which is an


important embedded filesystem. The reader is also introduced to the murky area of building a filesystem image from scratch, when what's available does not serve the purpose.

The chapter on MTD subsystems covers Memory Technology Devices, which is an abstraction layer between the low-level device complexities, and the higher-layer data organisation and storage formats that use memory devices.

Chapter 11 introduces the Swiss-army knife of embedded Linux development—BusyBox. This chapter, along with Chapters 12 and 13, talks about setting up an embedded development environment—with cross-compilers, profiling, tracing and debugging tools.

Chapter 14 is about kernel debugging techniques, and following it are debugging techniques for user-space application programs in Chapter 15.

In the next chapter, 'Porting Linux', the author takes a look at the various issues involved in porting Linux to a custom board. The book finally ends with real-time Linux in Chapter 17.

As I mentioned earlier, I would place this book in the reference section of my bookshelf, and pick it up whenever I am faced with one of the issues it deals with. The book would be of great help to embedded Linux developers who have some work experience. 

By: Amit Kumar Saha

The author has dabbled with the Linux kernel and pursues technical writing as a hobby. This means that whatever he writes here are his own views and not that of his employer's. He works for Sun Microsystems, India. His homepage is at [amitsaha.in.googlepages.com](http://amitsaha.in.googlepages.com)



# Envisage'09

## Festival To Celebrate Big Ideas



**E** (endeavour) N(navigate) V(vision) I(innovation) S(strife) A(analyse) G(graft) E(excel)! The acronym explains clearly what young minds with a vision aim at, while striving for excellence.

Envisage'09, a national level technical festival organised on February 28 and March 1, 2009, proved to be yet another milestone for the Institute of Informatics and Communication (IIC), South Campus, University of Delhi (DU), to bring together the youth, academia and industry. It saw participation from enthusiastic and motivated groups of people from all over India, including DU, the IITs and the NITs. There was an overwhelming response from 137 teams from across the country, comprising over 250 participants.

With "Think, Connect, Create and Share" as the motto, and open source as the underlying theme, the fourth edition of Envisage brought the concept of 'Inclusive Development' into focus.

### The journey begins

This two-day mega fest, comprising a plethora of events, was aimed at providing a platform to promote the spirit of ideas.

At the inauguration ceremony, eminent personalities from industry and academia shared a common platform. Professor Dinesh Singh, director, University of Delhi, South Campus, Dr Sanjeev Singh, faculty, IIC, and coordinator of Envisage'09, Dr MK Dass, teacher in charge, IIC, kept the audience enthralled.

After the inaugural formalities, a number of events were flagged off, C mantics (on C coding), Bugfix (the debugging challenge) and Object-o-mania (object-oriented programming) saw the clash of the programming champs, while 'What Lies Beneath' (a logic and reasoning contest) and 'No Ugly Duckling' (an open source challenge) laid emphasis on analytical and reasoning skills. Eureka, the project presentation, saw a gruelling battle as the short-listed candidates explained their ideas and innovations to a large

audience, under the microscopic scrutiny of the judges.

'Battlefield', the LAN-based gaming contest, witnessed an intense clash of wits, precision and teamwork, with gaming enthusiasts leaving no stone unturned in their hunt for supremacy. 'Junkyard Wars', the PC assembling contest, attracted a number of participants aiming to become hardware champions.


Day 2 kicked off with 'Termi-net-er' (a networking and socket programming contest) and the much-awaited IT and communication quiz, COMMIT, which saw quite a number of enthusiastic spectators cheering the participants on. 'Cobweb', the Web designing contest, saw Internet geeks design websites on pre-defined themes.

The theme this time was to develop a framework. While the day was filled with fun and frolic, the night involved sipping mugs of tea and munching on snacks while tackling various tough algorithmic problems based on the theme of *Shumdog Millionaire*.

'Ad-Vantage' (creating ads) and 'Die Hard' (a series of fun games) were the biggest crowd pullers at Envisage'09, filling the halls to capacity. The crowds roared with appreciation at some spectacular performances by

the contestants. 'ClickIt' (a digital photography contest) also attracted a number of participants, showcasing their imaginative skills.

### The journey continues...

Envisage has pioneered the idea of providing a platform for those in the open source universe in the country, to meet and influence each other. If the momentum set here is anything to go by, do fasten your seat belts, for Envisage continues to grow bigger and better, and the level of excitement will only increase in the years to come. 

**By: Shivani Mehrotra**

Shivani is the student coordinator at IIC, University of Delhi, South Campus.



▲ Waiting with bated breath for the results: The Closing Ceremony